Submitted to the Annals of Statistics

arXiv:1704.07016v1 [stat.ME] 24 Apr 2017

A NEW SVD APPROACH TO OPTIMAL TOPIC ESTIMATION

By Zheng Tracy Ke and Minzhe Wang

University of Chicago

In the probabilistic topic models, the quantity of interest—a low-rank matrix consisting of topic vectors—is hidden in the text corpus matrix, masked by noise, and Singular Value Decomposition (SVD) is a potentially useful tool for learning such a low-rank matrix. However, the connection between this low-rank matrix and the singular vectors of the text corpus matrix are usually complicated and hard to spell out, so how to use SVD for learning topic models faces challenges.

We overcome the challenge by revealing a surprising insight: there is a low-dimensional *simplex* structure which can be viewed as a bridge between the low-rank matrix of interest and the SVD of the text corpus matrix, and which allows us to conveniently reconstruct the former using the latter. Such an insight motivates a new SVD-based approach to learning topic models.

For asymptotic analysis, we show that under the popular probabilistic model (Hofmann, 1999), the convergence rate of the ℓ^1-error of our method matches that of the minimax lower bound, up to a multi-logarithmic term. In showing these results, we have derived new element-wise bounds on the singular vectors and several large-deviation bounds for weakly dependent multinomial data. Our results on the convergence rate and asymptotical minimaxity are new.

We have applied our method to two data sets, Associated Process (AP) and Statistics Literature Abstract (SLA), with encouraging results. In particular, there is a clear simplex structure associated with the SVD of the data matrices, which largely validates our discovery.

1. Introduction. In text mining, the problem of topic estimation is of interest in many application areas such as digital humanities, computational social science, e-commerce, and government science policy (Blei, 2012).

Consider a setting where we have n (text, say) documents. The documents share a common vocabulary of p words, and each of them discusses one or more of the K topics. Typically, n and p are large and K is relatively small. Table 1 presents two data sets of this kind, which we analyze in this paper.

MSC 2010 subject classifications: Primary 62H12, 62H25; secondary 62C20, 62P25.

Keywords and phrases: Asymptotic minimaxity, Ideal Simplex, k-means, mixed membership, multinomial distribution, nonnegative matrix factorization, random matrix theory, SCORE, sin-theta theorem

TABLE 1
Two data sets for topic estimation

Data sets	Vocabulary	Documents	Topics
Associated Press (AP)	10473 words	2246 news articles	"crime", "politics", "finance"
Statistical Literature Abstracts (SLA)	2934 words	3193 abstracts	"multiple testing", "variable selection" "experimental design", "bayes" "spectral analysis", "application"

We adopt the *probabilistic Latent Semantic Indexing (pLSI)* model (Hofmann, 1999) which is popular in this area. Suppose we observe a matrix $D \in \mathbb{R}^{p,n}$ (called the *text corpus* matrix in the literature), where for $1 \le i \le n$ and $1 \le j \le p$, $D(j,i) \in [0,1]$ is the observed fraction of word j in document i, $1 \le i \le n$ and $1 \le j \le p$. Write

$$D = D_0 + (D - D_0) = \text{"}signal\text{"} + \text{"}noise\text{"}, \tag{1}$$

where $D_0(j,i)$ is the expected frequency of word j in document i, and $D(j,i) - D_0(j,i)$ represents the observational variation.

In $pLSI$, we impose a low-rank structure on D_0. In detail, for $1 \le i \le n$, we assume that document i discusses each of the K topics with weights prescribed by a Probability Mass Function (PMF) $w_i \in \mathbb{R}^K$, where

$$w_i(k) = \text{expected weight that document } i \text{ puts on topic } k, \quad 1 \le k \le K.$$

Also, given that document i is discussing topic k, the expected frequency that word j appears in document i is $A_k(j)$, where $A_k \in \mathbb{R}^p$ is a PMF that does not depend on individual documents. Write $A = [A_1, A_2, \ldots, A_K]$ and $W = [w_1, w_2, \ldots, w_n]$. Recalling that $D_0(j,i)$ is the expected frequency of word j in document i, it is seen that

$$D_0(j,i) = \sum_{k=1}^{K} A_k(j) w_i(k), \qquad \text{or equivalently} \qquad D_0 = AW.$$

Combining this with (1) gives

$$D = AW + (D - D_0), \qquad \text{"}signal\text{"} = AW, \qquad \text{"}noise\text{"} = (D - D_0). \tag{2}$$

Our main interest is to use D to estimate the "topic matrix" A.

Latent Dirichlet Allocation (LDA) (Blei, Ng and Jordan, 2003) is a Bayesian approach to topic modeling. It imposes a Dirichlet prior on the columns of W, and estimates A by a variational EM algorithm. Despite its popularity, LDA is relatively slow computationally, especially when (n,p) are large. Also, the convergence of LDA is not carefully studied in the literature.

Since the matrix of interest, A, is contained in the low-rank signal matrix $D_0 = AW$, a standard frequentist response contains two steps.

- A Singular Value Decomposition (SVD) step that simultaneously reduces the dimension and noise.
- An "oracle approach" step for reconstructing the topic matrix A.

In the second step, we first study how to reconstruct A in the oracle case where the noise is removed from our model (so D_0 is known), using the first few singular vectors of D_0. We then extend the idea to the real case.

The "separable NMF" algorithm in Arora, Ge and Moitra (2012) can be viewed as an oracle approach but it does not use an SVD step for dimension and noise reduction. As a result, the estimates could be more noisy than expected and the rate of convergence derived in their paper might be non-optimal. See Section 2.3 for more discussion.

Papadimitriou et al. (2000) and Bansal, Bhattacharyya and Kannan (2014) use an SVD step, but without an oracle approach step. Their approaches are shown to behave well in some examples presented in their papers, but whether they work satisfactorily in more general cases remains unclear.

For successful estimation, we recognize that it is critical to use both an SVD step and a well-thought oracle approach. In particular, the connection between the quantity of interest—the topic matrix A—and the SVD is indirect and opaque, and how to elucidate such a connection is the key.

DEFINITION 1.1. *We call word j an anchor word*[1] *if row j of A has exactly one nonzero entry, and an anchor word for topic k if the nonzero entry locates at column k, $1 \le k \le K$.*

It is believed that for each of the K topics, there are a few anchor words. This is also supported by empirical evidence; see Section 1.3.

Our main contributions are four-fold:

- (*Discovery of a simplex structure*). We construct a $p \times (K-1)$ matrix using the first K left singular vectors of D and view each row as a point in $\mathbb{R}^{K-1}$. The rows generate a point cloud with the silhouette of a simplex, where each "anchor row" falls close to one of the vertices, and each "non-anchor row" falls close to an interior point of the simplex.
- (*A new SVD approach*). The simplex structure elucidates the connection between the topic matrix A and SVD and suggests a straightforward-way for constructing A; this gives rise to a new SVD approach.
- (*Optimal rate*). Using the ℓ^1-loss as the loss function, we derive the optimal rate of convergence. As far as we know, such a result is new.

[1]The term was introduced by Arora, Ge and Moitra (2012), in connection to the separable conditions for Nonnegative Matrix Factorization (Donoho and Stodden, 2004).

- (*Achievability*). We show that the proposed SVD approach is rate optimal. Recently, the rates for several procedures have been studied, but these rates are unfortunately not sharp; see Section 2.3 for details.

1.1. *SVD and the Ideal Simplex (oracle case).* Consider the matrix consisting of the first K left singular vectors of D. The matrix provides a good estimate of $col(A)$,[2] but not the matrix A itself: two matrices are related to each other by an unknown $K \times K$ transformation matrix.

How to estimate this matrix poses challenges. We approach this by relating it to the vertices of a *simplex* in $\mathbb{R}^{K-1}$, where the latter can be conveniently estimated. To spell out the idea, we start with the *oracle* case (where D_0 is assumed as known). The *real* case is studied in Section 1.2.

When processing the documents, people often re-normalize the rows of D (in the oracle case, the matrix D_0) to address the issue that some words appear much less frequently than the others. In light of this, introduce

$$M_0 = \mathrm{diag}(n^{-1}D_0 1_n),^{3} \qquad (1_n\text{: } n\text{-dimensional vector of 1's}),$$

and consider the normalized (oracle) data matrix $M_0^{-1/2}D_0$.[4]

Let $\sigma_1 > \sigma_2 > \ldots > \sigma_K$ be the first K singular values of $M_0^{-1/2}D_0$, and let $\xi_1, \xi_2, \ldots, \xi_K$ be the corresponding left singular vectors. Write $\Xi = [\xi_1, \xi_2, \ldots, \xi_K]$. Since $D_0 = AW$, $col(M_0^{-1/2}A) = col(\Xi)$ (see footnote 2), so there is a transformation matrix $V \in \mathbb{R}^{K,K}$ (compare above) such that

$$\Xi = M_0^{-1/2}AV.$$

Note that in the oracle SVD approach, the goal is to use Ξ to estimate A (assuming M_0 is given), and to this end, the key is to use Ξ to estimate V.

Write $V = [V_1, V_2, \ldots, V_K] = [v_1, v_2, \ldots, v_K]'$, such that V_k and v_k' are the k-th column and row, respectively. Define two *matrices of entry-wise ratios* $R \in \mathbb{R}^{p,K-1}$ and $V^* \in \mathbb{R}^{K,K-1}$ by

$$R(j,k) = \xi_{k+1}(j)/\xi_1(j), \qquad 1 \le j \le p,\ 1 \le k \le K-1,$$

and

$$V^*(\ell,k) = V_{k+1}(\ell)/V_1(\ell), \qquad 1 \le \ell \le K,\ 1 \le k \le K-1.$$

Here R is obtained by taking the ratio between each of $\xi_2, \ldots, \xi_K$ and ξ_1 in an entry-wise fashion, V^* is obtained from $V_1, \ldots, V_K$ similarly.

[2]The notation $col(A)$ stands for the linear space spanned by the columns of matrix A.

[3]For a vector $d \in \mathbb{R}^n$, $\mathrm{diag}(d)$ denotes the $n \times n$ diagonal matrix whose i-th diagonal entry is the i-th entry of d, $1 \le i \le n$.

[4]Such a normalization is not necessary for the oracle case but is critical in the real case for controlling the errors introduced by noise.

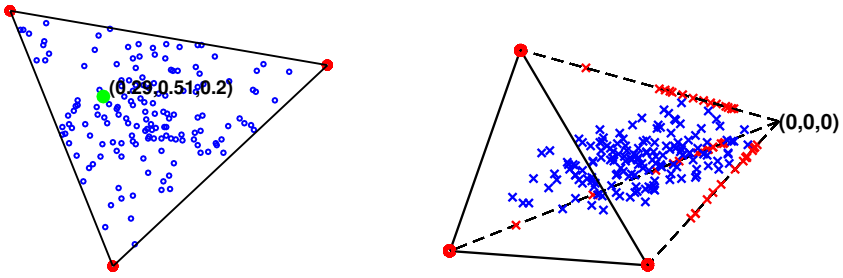

FIG 1. $K = 3$. *Left panel: Ideal Simplex (solid triangle). Each circle represents a row of R (red: anchor words, blue: non-anchor words). Every r_j is a convex combination of the K vertices, where the weight for one r_j is displayed. Right panel: Why it is appropriate to use entry-wise eigen-ratios. The solid triangle is the simplex formed by rows of $\tilde{A}V$. Each cross represents a row of Ξ; these rows are obtained by rescaling the rows of $\tilde{A}V$, so they no longer have the silhouette of a simplex.*

Write $R = [r_1, r_2, \ldots, r_p]'$ and $V^* = [v_1^*, v_2^*, \ldots, v_K^*]'$ so that $r_j' \in \mathbb{R}^{K-1}$ is the j-th row of R and $(v_k^*)' \in \mathbb{R}^{K-1}$ is the k-th row of V^*. The following lemma is one of our key observations.[5]

LEMMA 1.1 (Ideal Simplex). *The rows of R form a point cloud with the silhouette of a simplex $\mathcal{S}_K^*$ with $v_1^*, v_2^*, \ldots, v_K^*$ being the vertices.*

- *If word j is an anchor word, then r_j falls on one of the vertices of $\mathcal{S}_K^*$.*
- *If word j is a non-anchor word, then r_j falls into the interior of $\mathcal{S}_K^*$ (or the interior of an edge/face), and equals to a convex combination of $v_1^*, v_2^*, \ldots, v_K^*$ with π_j being the weight vector, where π_j' is the j-th row of the matrix $\Pi = [\mathrm{diag}(\xi_1)]^{-1} M_0^{-1/2} \cdot A \cdot \mathrm{diag}(V_1)$, $1 \le j \le p$.*

We can now use (M_0, ξ_1, R) to reconstruct $A = [A_1, A_2, \ldots, A_K]$.

- (*Vertices Hunting*). Use rows of R and the simplex structure to locate all vertices $v_1^*, v_2^*, \ldots, v_K^*$.
- (*Weight matrix reconstruction*). Use r_j and the vertices to obtain π_j, $1 \le j \le p$. This gives us the matrix Π.
- (*Topic matrix reconstruction*). For each $1 \le k \le K$, A_k can be easily reconstructed from that A_k is a PMF and that A_k is proportional to the k-th column of $(M_0^{1/2} \cdot \mathrm{diag}(\xi_1) \cdot \Pi)$ (as dictated by the lemma).

As far as we know, our approach is new. Also, the simplex structure is low-dimensional and is based on the entry-wise ratios, and is very different from

[5]This lemma holds for any diagonal matrix M_0 that is positive-definite.

those in Donoho and Stodden (2004); Arora, Ge and Moitra (2012) (which are high dimensional). See Figure 1 (left panel).

Remark. We explain the rationale of taking entry-wise eigen-ratios. Write $A = [A_1, A_2, \ldots, A_K] = [a_1, a_2, \ldots, a_p]'$ and let $\tilde{A} = [\tilde{a}_1, \tilde{a}_2, \ldots, \tilde{a}_p]'$, where $\tilde{a}_j = h_j^{-1} a_j$ with $h_j = \|a_j\|_1$. Recalling $\Xi = M_0^{-1/2} AV$, it follows that

$$\Xi = M_0^{-1/2}\mathrm{diag}(h) \cdot \tilde{A}V, \qquad \text{(note each row of } \tilde{A} \text{ has a unit } \ell^1\text{-norm)}.$$

The rows of Ξ represent p points in $\mathbb{R}^K$ that fall in the simplicial cone spanned by $\{0\}$ and $\{v_1, \ldots, v_K\}$. See Figure 1 (right panel). How to identify $v_1, v_2, \ldots, v_K$, with which the cone is generated, is nontrivial, and the problem is even harder when we only observe $\hat{\Xi}$, not Ξ. To solve the problem, one possibility is to estimate the diagonal matrix $M_0^{-1/2}\mathrm{diag}(h)$ and then use it to estimate $\tilde{A}V$; once we have a good estimate of $\tilde{A}V$, we can estimate V similarly by vertices hunting (the rows of $\tilde{A}V$ represent p points in $\mathbb{R}^K$ where an anchor row falls on a vertex of the simplex $\mathcal{S}(v_1, v_2, \ldots, v_K)$, and a non-anchor row falls in the interior). However, it is unclear how to estimate $M_0^{-1/2}\mathrm{diag}(h)$ (note that h depends on the unknown A). *Our idea is different. We recognize that this diagonal matrix is largely nuisance: Its effect can be conveniently removed by taking entry-wise ratios as proposed so there is no need to estimate it directly. In fact, it is seen that the entry-wise eigen-ratio matrix R equals to ΠV^* and does not depend on $M_0^{-1/2}\mathrm{diag}(h)$: the diagonals of the latter are completely cancelled out when we take entry-wise ratios!* An alternative to cancel out these diagonals is to normalize each row of Ξ to have an unit ℓ^q-norm for some $q > 0$. But when we do this, the geometry associated with the resultant matrix is more complicated, for each of its rows falls on the surface of the unit ℓ^q ball. This makes the problem unnecessarily more complicated.

Remark. The idea of using entry-wise eigen-ratios for inference was introduced earlier in Jin (2015) and Jin, Ke and Luo (2016), but in a very different setting (i.e., analysis of social networks). Their models are very different from ours, both in mathematical forms and in statistical distributions, and the quantities of interest are also very different. We have to derive the entry-wise ratio matrix and the simplex structure by ourselves, and the statistical analysis in those papers is also very different from ours.

1.2. *A novel SVD approach to topic estimation (real case).* In the real case, we only observe a "blurred" version of the matrix R and so a "blurred" version of the Ideal Simplex. It is therefore challenging to have a computationally feasible approach to estimate the vertices of the Ideal Simplex.

One possible approach is to use optimization: for an objective function of a set of K vertices, we optimize it over all possible sets of vertices. However, it is unclear how to choose the objective function for our purpose; distance from the farthest data point to the simplex is a handy choice, but we have to penalize for the cases of extreme volumes, and it is unclear how to choose the penalty appropriately. Also, solving for the optimizer may be time consuming, especially if the objective function is non-linear and non-convex.

We propose an approach that is fast and easy-to-use. Let $\hat{\sigma}_1 > \hat{\sigma}_2 > \ldots > \hat{\sigma}_K$ be the K leading singular values of $M^{-1/2}D$, where $M = \mathrm{diag}(\frac{1}{n}D1_n)$, and let $\hat{\xi}_1, \hat{\xi}_2, \ldots, \hat{\xi}_K$ be the corresponding left singular vectors. Fixing a threshold $t > 0$, we let $\hat{R}$ be the (regularized) empirical counterpart of R, where for $1 \le k \le K-1, 1 \le j \le p$,

$$\hat{R}(j,k) = \begin{cases} \hat{\xi}_{k+1}(j)/\hat{\xi}_1(j), & \text{if } |\hat{\xi}_{k+1}(j)/\hat{\xi}_1(j)| \le t, \\ t, & \text{if } \hat{\xi}_{k+1}(j)/\hat{\xi}_1(j) > t, \\ -t, & \text{if } \hat{\xi}_{k+1}(j)/\hat{\xi}_1(j) < -t. \end{cases} \tag{3}$$

Write $\hat{R} = [\hat{r}_1, \hat{r}_2, \ldots, \hat{r}_p]'$ so that $\hat{r}_j'$ is the j-th row of $\hat{R}$. We view each $\hat{r}_j$ as a point in $\mathbb{R}^{K-1}$. Our idea is to first cluster $\hat{r}_1, \ldots, \hat{r}_p$ into L clusters, where L is a tuning integer that is usually a few times larger than K. We compute the center for each cluster, in hopes that some of the centers fall close to the vertices of the Ideal Simplex. We then select K centers (out of L of them) as the estimates of the vertices of the ideal simplex.

In detail, for any affinely independent vectors $a_1, a_2, \ldots, a_K \in \mathbb{R}^{K-1}$, let $\mathcal{S}(a_1, a_2, \ldots, a_K)$ be the simplex with vertices $a_1, a_2, \ldots, a_K$. For any $b \in \mathbb{R}^{K-1}$, let $\mathrm{distance}(b, \mathcal{S}(a_1, a_2, \ldots, a_K))$ denote the Euclidean distance between b and $\mathcal{S}(a_1, a_2, \ldots, a_K)$ if b falls outside of the simplex, and 0 otherwise (the distance can be computed conveniently via quadratic programming). We propose the following two-stage *Vertices Hunting (VH)* algorithm. Input: K, a tuning integer $L > K$, and $\hat{r}_1, \cdots, \hat{r}_p$. Output: estimated vertices $\hat{v}_1^*, \cdots, \hat{v}_K^*$ (see Figure 2 for illustration).

VH-1. Cluster by applying the classical k-means to $\hat{r}_1, \cdots, \hat{r}_p$, assuming there are L clusters. Let $\hat{\theta}_1, \cdots, \hat{\theta}_L$ be the Euclidean centers of the clusters.

VH-2. Let $1 \le \hat{j}_1 < \hat{j}_2 < \cdots < \hat{j}_K \le L$ be the indices such that $\hat{\theta}_{\hat{j}_1}, \cdots, \hat{\theta}_{\hat{j}_K}$ are affinely independent and minimize

$$\max_{1 \le j \le L} \big\{ \mathrm{distance}\big(\hat{\theta}_j,\ \mathcal{S}(\hat{\theta}_{j_1}, \cdots, \hat{\theta}_{j_K})\big) \big\}.$$

Output $\hat{v}_k^* = \hat{\theta}_{\hat{j}_k}$, $1 \le k \le K$. If no such $(\hat{j}_1, \cdots, \hat{j}_K)$ exist, output $\hat{v}_1^* = (0, \ldots, 0)'$ and $\hat{v}_{k+1}^* =$ the k-th standard basis vector of $\mathbb{R}^{K-1}$.

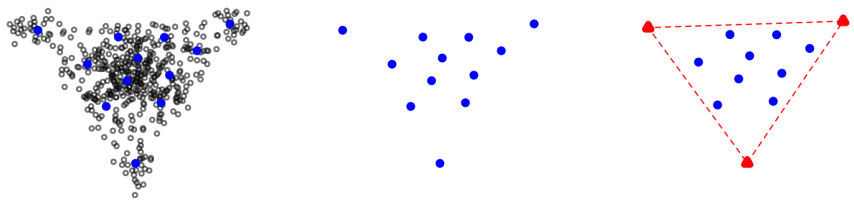

FIG 2. *Vertices Hunting algorithm ($K = 3$). Left: Apply the classical k-means to $\hat{r}_1, \ldots, \hat{r}_p$ and obtain the Euclidean centers of clusters (blue points). Middle: Remove $\hat{r}_1, \ldots, \hat{r}_p$ and only keep the cluster centers. Right: Fit a simplex using these cluster centers.*

We propose the following topic estimation method, which starts with the vertices hunting algorithm and extends the gained insight in the oracle case. Input: D, K, a tuning threshold $t > 0$ and a tuning integer $L > K$. Output: $\hat{A}$, an estimate of A.

1. (*Vertices Hunting*). Obtain $\hat{R} = [\hat{r}_1, \hat{r}_2, \ldots, \hat{r}_p]'$ as in (3), and obtain $\hat{v}_1^*, \cdots, \hat{v}_K^*$ by applying the Vertices Hunting algorithm above.
2. (*Weight matrix estimation*). For $1 \leq j \leq p$, solve $\hat{\pi}_j^*$ from

$$\begin{pmatrix} 1 & \ldots & 1 \\ \hat{v}_1^* & \ldots & \hat{v}_K^* \end{pmatrix} \hat{\pi}_j^* = \begin{pmatrix} 1 \\ \hat{r}_j \end{pmatrix}.$$

   Regularize $\hat{\pi}_j^*$ by setting all negative entries to 0 and then re-normalizing the vector to have a unit sum; denote the resultant vector by $\hat{\pi}_j$. Write $\hat{\Pi} = [\hat{\pi}_1, \cdots, \hat{\pi}_p]'$.
3. (*Topic matrix estimation*). Form the matrix $\hat{A}^* = M^{1/2} \cdot \mathrm{diag}(\hat{\xi}_1) \cdot \hat{\Pi}$, where $M = \mathrm{diag}(\frac{1}{n} D 1_n)$. Normalize each column of $\hat{A}^*$ to have a unit ℓ^1-norm. The resultant matrix is our output matrix $\hat{A}$.

We have two tuning parameters: L (number of clusters) and t (threshold in (3)). We recommend $L = 10K$. We also recommend $t = 2\log(n)$ for large (n, p), but in our data analysis (simulated and real), thresholding in (3) does not have a major effect, so we take $t = \infty$ for simplicity.

The computing cost of our procedure has three parts: SVD, local center estimation by k-means, and vertices estimation by exhaustive search. The SVD is a rather manageable algorithm even for large matrices, and the k-means (despite that its comparably high complexity, it is not NP hard [6]) is usually executed in practice by the Llyod algorithm, which is pretty fast. The

[6]We may have the wrong impression that the k-mean is always NP-hard: the k-means is NP-hard if both the dimension and the number of clusters are large, but this is not the case here for both of them (namely, $(K-1)$ and L) are reasonably small.

exhaustive search could be relatively slow when both (K, L) are large (and is reasonably fast otherwise), but since it aims to solve a simple problem, it can be replaced by some much faster greedy algorithm. How to improve this part is not the main focus of the paper, so we leave it to the future work.

Remark. Our procedure is very flexible and the main idea continues to work if we revise some steps. For example, the method continues to work if we use a different normalization matrix M (see footnote 5), or replace the k-means by some other clustering algorithms (e.g., k-median or an $(1+\epsilon)$-approximate solution of k-means). Also, if we know which are the anchor words (say, by prior knowledge or by some anchor-selection algorithms), we can revise our algorithm accordingly to accommodate such a situation.

1.3. *Real data applications.* We implement our method to the two data sets in Table 1.

Associated Press (AP) data. The AP data set (Harman, 1993) consists of 2246 news articles with a vocabulary of 10473 words. For preprocessing, we remove 191 stop words, and keep the 8000 words that appear most frequently in the vocabulary. We also remove 5% of the documents that are among the shortest.

How to determine the number of topics K is a challenging problem. The scree plot suggests $K = 3$, and we have applied our method with $K = 2, 3, \ldots, 6$ and it seems $K = 3$ gives the most reasonable results.

We now report some results for $K = 3$. First, Table 2 presents the top 15 representative words for the each of the three topics in (a word is called "representative" of a topic if its corresponding $\hat{r}_i$ is close to the estimated vertex of that topic). The results suggest that the three estimated topics can be interpreted as "crime", "politics", and "finance", respectively.

TABLE 2
Top 15 representative words for each estimated topic in the AP data ($K = 3$).

"Crime"	*shootings, injury, mafia, detective, bangladesh, dog, hindus, gunfire, aftershocks, bears, accidentally, handgun, unfortunate, dhaka, police*
"Politics"	*eventual, gorbachevs, openly, soviet, primaries, sununu, yeltsin, cambodia, torture, soviets, herbert, gephardt, afghanistan, citizenship, popov*
"Finance"	*trading, stock, edged, dow, rose, traders, stocks, indicators, exchange, share, guilders, bullion, lire, christies, unleaded*

Also, Figure 3 plots the rows of the matrix $\hat{R}$ (see (3)). Since $K = 3$, each row or $\hat{R}$ is a point in $\mathbb{R}^2$. The data cloud illustrates the silhouette of a triangle, which fits very well with our theory on the simplex structure.

In Figure 3, it is interesting to note that there is a "hole" near the edge connecting the two vertices of "crime" and "finance." This makes perfect

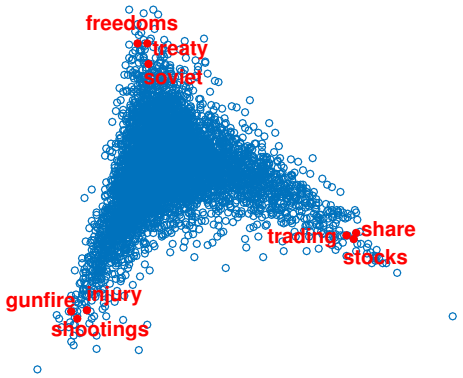

FIG 3. *Plot of $\hat{r}_1, \ldots, \hat{r}_p$ (data: Associated Press; $K = 3$). A triangle is visible in the data cloud, where the three vertices represent the three topics "crime", "politics", and "finance". Red: identified nearly-anchor words.*

sense: words that are related to both "crime" and "finance" tend to be also related to "politics". In contrast, there are many words that are related to both "politics" and "crime" but are unrelated to "finance", and there are many words that are related to both "politics" and "crime" but are unrelated to "crime", so we do not see a hole near either of the other two edges. We thank Art Owen for very helpful comments.

Statistical Literature Abstracts (SLA) data. This data set was collected by Ji and Jin (2016) (see also Kolar and Taddy). It consists of the abstracts of 3193 papers published in *Annals of Statistics*, *Biometrika*, *Journal of the American Statistical Association*, and *Journal of the Royal Statistical Society: Series B*, from 2003 to the first half of 2012. The full vocabulary contains 2934 words. For preprocessing, we remove 209 stop words. We also remove 40% of the documents that are among the shortest.

We tried our method with $K = 2, 3, \ldots, 6, 7, 8$ and found that $K = 6$ yields the most meaningful results, so we pick $K = 6$ for our study. Table 3 shows the top 15 representative words in each of the six estimated topics. These topics can be interpreted as "Multiple Testing", "Bayes", "Variable Selection", "Experimental Design", "Spectral Analysis", and "Application".

1.4. *Connections.* We propose an SVD-based geometrical approach to topic estimation, at the heart of which is the discovery of a low-dimensional simplex that is associated with the matrix of entry-wise eigen-ratios.

Ordinary SVD taught in textbooks is useful for both dimension reduction and noise reduction. However, for many modern applications (e.g., cancer clustering, network community detection), ordinary SVD is frequently found

TABLE 3
Top 15 representative words for each estimated topic in the SLA data ($K = 6$).

"Multiple Testing"	*stepup, stepdown, rejections, hochberg, fwer, singlestep, familywise, benjamini, bonferroni, simes, intersection, false, rejection, positively, kfwer*
"Bayes"	*posterior, prior, slice, default, credible, conjugate, priors, improper, wishart, admissible, sampler, tractable, probit, normalizing, mode*
"Variable Selection"	*angle, penalties, zeros, sure, selector, selection, stability, enjoys, penalization, regularization, lasso, tuning, irrelevant, selects, clipped*
"Experimental Design"	*aberration, hypercube, latin, nonregular, spacefilling, universally, twofactor, blocked, twolevel, designs, crossover, resolution, factorial, toxicity, balanced*
"Spectral Analysis"	*trajectories, amplitude, eigenfunctions, realizations, away, gradient, spectra, discrimination, functional, auction, nonstationarity, spacetime, slex, curves, jumps*
"Application"	*instrument, vaccine, instruments, severity, affects, compliance, infected, depression, schools, assignment, participants, causal, warming, rubin, randomized*

to be unsatisfactory. To better use such a powerful tool, many improved SVD approaches are proposed. These include the sparse PCA approach (Zou, Hastie and Tibshirani, 2006; Johnstone and Lu, 2009) and the IF-PCA approach (Jin and Wang, 2016). Our approach is also an improved SVD approach, and at a high level, it is connected to these works aforementioned, but of course it is also very different.

Our approach is also connected to other geometrical approaches to topic estimation. In a related setting, Donoho and Stodden (2004) pointed out that the rows of D_0, viewed as points in $\mathbb{R}^n$, live in a simplicial cone with K supporting rays, and if we normalization each row of D_0 by its ℓ^1-norm, the point cloud forms a simplex; such a geometrical structure was used in Ding et al. (2013) for topic estimation. Arora, Ge and Moitra (2012) utilized the geometrical structure in the "word-word co-occurence" matrix $Q = DD'$ for topic estimation. In the oracle case (where $D = D_0$), the rows of Q, under a proper normalization, form a simplex in $\mathbb{R}^p$ and has K vertices. We compare these simplex structures with ours in Table 4.

TABLE 4
Comparison of the simplex structures generated from different source matrices.

Source	Oracle counterpart	Normalize by	Dimension
text corpus D	$AW(= D_0)$	row-wise ℓ_1-norm	n
word co-occurrence Q	$AWW'A'$	row-wise ℓ_1-norm	p
singular vectors $\hat{\Xi}$	$AV(= \Xi)$	first column	$K-1$

The simplex we discover is SVD-based so both dimension reduction and noise reduction are implied, and is very different in nature.

- Our simplex lives in $\mathbb{R}^{K-1}$ while the other simplices live in $\mathbb{R}^n$ or $\mathbb{R}^p$ (in practice, $\min\{n, p\}$ is usually a few thousands while K is usually small (e.g., $K \leq 20$)). A low-dimensional simplex structure makes vertices

hunting both statistically more accurate and computationally faster.
- Our simplex is built on $\hat{\Xi}$ while the other simplices are built on D or Q, where due to the noise-removal effect of SVD, $\hat{\Xi}$ is much less noisy than D or Q. This also makes our approach statistically more accurate than the other approaches aforementioned.

These intuitions are confirmed by our asymptotic analysis. In Section 2, we show that our method achieves the optimal rate of $(Nn)^{-1/2}\sqrt{p}$, while the error rate of the method in Arora, Ge and Moitra (2012) is between $(Nn)^{-1/2}p$ to $(Nn)^{-1/2}p^4$ (which is much slower and thus non-optimal).

1.5. *Content and notations.* The remaining part of this paper is organized as follows: Section 2 states the main results (rate of convergence and optimality). Section 3 develops the key technical tools and proves the rate of convergence. Section 4 contains numerical experiments, and Section 5 contains discussions. The proofs are relegated to Sections 6-7.

Through out this paper, $\mathbb{R}$ denotes the set of real numbers, $\mathbb{R}^p$ denotes the p-dimensional real Euclidean space, and $\mathbb{R}^{p,q}$ denotes the set of $p \times q$ real matrices. For two positive sequences $\{a_n\}_{n=1}^{\infty}$ and $\{b_n\}_{n=1}^{\infty}$, we write $a_n = O(b_n)$, $a_n = o(b_n)$, and $a_n \lesssim b_n$, if $\lim_{n\to\infty}(a_n/b_n) < \infty$, $\lim_{n\to\infty}(a_n/b_n) = 0$, and $\limsup_{n\to\infty}(a_n/b_n) \leq 1$, respectively. Given $0 \leq q \leq \infty$, for any vector x, $\|x\|_q$ denotes the L_q-norm of x. For any matrix M, $\|M\|$ denotes the spectral norm of M and $\|M\|_F$ denotes the Frobenius norm of M. When M is symmetric, $\lambda_{\max}(M)$ and $\lambda_{\min}(M)$ denote the maximum and minimum eigenvalues of M, respectively.

2. Main results. We describe the model and regularity conditions in Section 2.1 and state the main results in Section 2.2; our main results include the convergence rate of our method and its optimality. Section 2.3 contains a few remarks and compares our theoretical results with existing ones.

2.1. *Model and regularity conditions.* Let $A = [A_1, \ldots, A_K]$ and $W = [w_1, \ldots, w_n]$ be the same as in Section 1. For simplicity, we assume all documents have the same length N. The *pLSI* model (Hofmann, 1999) is often described as follows: The documents are generated independently. Let x_{im} be the m-th word of document i, $1 \leq i \leq n$, $1 \leq m \leq N$. For $i = 1, \ldots, n$, the words $\{x_{im}\}_{m=1}^{N}$ are drawn independently and identically, where for each $1 \leq m \leq N$, a topic $y_{im} \in \{1, 2, \ldots, K\}$ is first drawn using the PMF w_i and then x_{im} is drawn using the PMF $A_{y_{im}}$. According to this model, for each word $1 \leq j \leq p$ in the vocabulary, its expected frequency in document i is $\sum_{k=1}^{K} w_i(k)A_k(j)$. This implies $D_0 = AW$, as dictated in (2).

Write $D = [d_1, d_2, \ldots, d_n]$. We rewrite the *pLSI* model equivalently as

$$d_1, \ldots, d_n \text{ are independent}, \quad Nd_i \sim \text{Multinomial}(N, d_i^0), \quad 1 \le i \le n, \tag{4}$$

where

$$D_0 = [d_1^0, d_2^0, \ldots, d_n^0] = AW. \tag{5}$$

We adopt an asymptotic framework where K is fixed and (n, N, p) tend to infinity. Denote by a_j' the j-th row of A, $1 \le j \le p$. Introduce

$$H = \text{diag}(h_1, \ldots, h_p), \qquad \text{where } h_j = \|a_j\|_1, \text{ for } 1 \le j \le p.$$

DEFINITION 2.1. *$\Sigma_A = A'H^{-1}A$ is called the "topic-topic correlation" matrix, and $\Sigma_W = n^{-1}WW'$ is called the "topic-topic concurrence" matrix.*

Both Σ_A and Σ_W are properly scaled, because it always holds that $\lambda_{\max}(\Sigma_A) \le 1$, $\text{trace}(\Sigma_A) \ge 1$, $\lambda_{\max}(\Sigma_W) \le 1$ and $\text{trace}(\Sigma_W) \ge K^{-1}$. We assume for a constant $c_1 \in (0, 1)$,

$$\lambda_{\min}(\Sigma_W) \ge c_1, \qquad \lambda_{\min}(\Sigma_A) \ge c_1, \qquad \min_{1 \le k, \ell \le K} \Sigma_A(k, \ell) \ge c_1. \tag{6}$$

The last inequality of (6) is a realistic assumption, noting that in many real datasets all the topics share a considerable fraction of common words.

Let $M_0 = \text{diag}(\frac{1}{n} D_0 1_n)$ be as in Section 1.2. Introduce the $K \times K$ matrix

$$\widetilde{\Sigma}_A = A' M_0^{-1} A.$$

This matrix is also properly scaled (see the proof of Lemma 6.1). Unlike Σ_A which is determined by the topic matrix A only, $\widetilde{\Sigma}_A$ depends on both A and W. We assume for a constant $c_2 > 0$,

$$Eiggap\big(\Sigma_W \widetilde{\Sigma}_A\big) \ge c_2, \tag{7}$$

where $Eiggap(\cdot)$ denotes the minimum gap between any two singular values.

Denote by $\mathcal{T}_k$ the set of anchor words of topic k, for $1 \le k \le K$. Let h_j's be the same as above. We assume for sequences $\{m_p\}_{p=1}^{\infty}$ and $\{\delta_p\}_{p=1}^{\infty}$,

$$\min_{1 \le k \le K} |\mathcal{T}_k| \ge m_p, \qquad \min_{j \in (\cup_{k=1}^K \mathcal{T}_k)} h_j \ge \delta_p. \tag{8}$$

The parameter δ_p connects to the definition of "separability" for topic matrices (Arora, Ge and Moitra, 2012): by (8), the topic matrix is δ_p-separable.

Let $\mathcal{C} = (\cup_{k=1}^K \mathcal{T}_k)^c$ denote the set of non-anchor words. For $1 \le j \le p$, let $\tilde{a}_j = h_j^{-1} a_j$, where a_j and h_j are the same as above; note that $\tilde{a}_j$ is a weight vector. For any integer $L \ge 1$, the k-means loss associated with the $\tilde{a}_j$'s of non-anchor words is defined as

$$RSS_n(L) = \min_{\eta_1^*,\ldots,\eta_L^*} \sum_{j\in\mathcal{C}} \big\{ \min_{1\le \ell\le L} \|\tilde{a}_j - \eta_\ell^*\|^2 \big\}.$$

Let $e_1, \ldots, e_K$ be the standard basis vectors of $\mathbb{R}^K$. We assume for a constant $c_3 > 0$ and a finite integer L_0,

$$\min_{j\in\mathcal{C}} \min_{1\le k\le K} \|\tilde{a}_j - e_k\| \ge c_3, \qquad RSS_n(L_0) \le \frac{m_p}{\log(n)}. \tag{9}$$

This assumption requires that the $\tilde{a}_j$'s of non-anchor words have mild "concentration." It is mainly for the convenience of analyzing the vertices hunting algorithm and can be largely relaxed.

2.2. *The rate of convergence (upper bound and lower bound).* Since the topic matrix A is self-normalized, it is appropriate to measure the loss using the ℓ_1 estimation error:

$$\mathcal{L}(\hat{A}, A) \equiv \min_{\kappa\text{: a permutation on } \{1,\ldots,K\}} \Big\{ \sum_{k=1}^K \|\hat{A}_k - A_{\kappa(k)}\|_1 \Big\}. \tag{10}$$

On top of this definition, we further "allocate" the error to each individual word. Recall that $\hat{a}_j'$ and a_j' are the respective j-th row of $\hat{A}$ and A, $1 \le j \le p$. Let $\mathcal{P}_K$ be the set of all $K \times K$ permutation matrices (i.e., each row/column has exactly one nonzero entry which is equal to 1). It is seen that

$$\mathcal{L}(\hat{A}, A) = \min_{T\in\mathcal{P}_K} \Big\{ \sum_{j=1}^p \|T\hat{a}_j - a_j\|_1 \Big\}.$$

We provide not only a bound for $\mathcal{L}(\hat{A}, A)$ but also a bound for each individual $\|T\hat{a}_j - a_j\|_1$. The following theorem is proved in Section 3.

THEOREM 2.1 (Upper bound). *Consider models (4)-(5), where (6)-(9) hold, $\min\{N, p\} \ge \log(n)$, $p\log^2(n)/(Nn) \le 1$, and $p \le CN^{3/4}$. Additionally, suppose $h_{\min} \ge C^{-1}/p$ and $m_p \ge p^2\log^2(n)/(Nn)$. In our method, set $t = \infty$ and $L \ge L_0 + K$, where L_0 is as in (9). With probability $1 - o(n^{-3})$, there exists a permutation matrix $T \in \mathcal{P}_K$ such that*

$$\frac{\|T\hat{a}_j - a_j\|_1}{\|a_j\|_1} \le C\sqrt{\frac{p\log(n)}{Nn}}, \qquad \textit{for all } 1 \le j \le p.$$

Moreover, with probability $1 - o(n^{-3})$,

$$\mathcal{L}(\hat{A}, A) \leq C\sqrt{\frac{p\log(n)}{Nn}}.$$

Let $\mathcal{A}_p(K)$ be the set of all $p \times K$ topic matrices $A = [a_1, \ldots, a_p]'$ (i.e., each column of A is a weight vector). Let $\mathcal{W}_n(K)$ be the set of all matrices $W \in \mathbb{R}^{n,K}$ such that each column of W is a weight vector. For each constant $c \in (0, 1]$, define

$$\Phi_{n,N,p}(K, c) = \Big\{(A, W) : A \in \mathcal{A}_p(K), W \in \mathcal{W}_n(K), \min_{1\leq j\leq p} \|a_j\|_1 \geq cp^{-1}\Big\}.$$

The following theorem is proved in Section 6.

THEOREM 2.2 (Lower bound). *Consider models* (4)-(5), *where* $p = o(Nn)$ *as* (n, N, p) *tend to infinity. For each* $c \in (0, 1]$, *there exist constants* $C_0 > 0$ *and* $\delta_0 \in (0, 1)$ *such that*

$$\inf_{\hat{A}} \sup_{(A,W)\in\Phi_{n,N,p}(K,c)} \mathbb{P}\left(\mathcal{L}(\hat{A}, A) \geq C_0\sqrt{\frac{p}{Nn}}\right) \geq \delta_0.$$

Combining the above two theorems, we conclude that our method achieves the optimal rate of convergence up to a logarithmic factor.

2.3. *Connections.* Recall that δ_p is the separability parameter and m_p is the minimum number of anchor words across different topics. Our method only requires very mild assumptions on (δ_p, m_p): $\delta_p \geq cp^{-1}$ and $m_p \geq p \cdot [p\log^2(n)/(Nn)]$ (the term in the bracket is very small). Somewhat surprisingly, as long as the two assumptions hold, neither δ_p nor m_p appears in the error rate. In fact, (δ_p, m_p) mainly affect the vertices hunting step, where the error rate is not the dominating term. Arora, Ge and Moitra (2012) is the first work that provides an explicit error rate for topic model estimation, and their results are still used as a benchmark by many literatures. However, the ℓ_1-estimation error rate of their method is $O((Nn)^{-1/2}(\delta_p^{-3}p))$, which is comparably slower than ours, especially when δ_p is small (for example, when δ_p ranges between $O(1/p)$ and $O(1)$, their rate falls between $O((Nn)^{-1/2}p)$ and $O((Nn)^{-1/2}p^4)$). Also, their error rate depends explicitly on δ_p, while neither the minimax rate nor the rate of our method depends on δ_p. Note that in many applications, δ_p tends to be rather small.

Recall that h_j is the average frequency of word j across different topics. In our upper/lower bound arguments, we assume $\min_{1\leq j\leq p} h_j \geq cp^{-1}$.

This assumption is quite reasonable from a practical perspective: In data preprocessing, a list of most infrequent words are often removed from the vocabulary. The assumption can also be relaxed. For example, we can extend our upper/lower bound arguments to a sparse setting, where no more than s_p ($s_p \ll p$) words have a nonzero h_j and each of the nonzero h_j is no less than $O(s_p^{-1})$. We conjecture that the minimax rate is $(Nn)^{-1/2}\sqrt{s_p}$ (up to a logarithmic factor) and our method is rate optimal provided we add an additional screening step to our procedure.

We have also assumed $p \leq CN^{3/4}$ but this is mostly for convenience of presentation. The analysis can be easily extended to the case of $p \gg CN^{3/4}$, except that the rate will be more complicated (see Theorem 3.1 for details).

Our result is connected to the minimax results for classical PCA (Birnbaum et al., 2013), but is also very different. Our result can not be deduced from existing results for PCA: the noise distribution in our setting is very different from those in the literature, and the goal here is also very different, where PCA is only one part of our procedure.

3. Proof of the upper bound. In this section, we prove Theorem 2.1. The proof requires two key results: (i) large-deviation bounds of the empirical singular vectors, and (ii) characterization of the noise accumulation in our method; these are addressed in Sections 3.1-3.2, respectively. In the proof, we have developed new technical tools, including an element-wise perturbation bound for eigenvectors (Lemma 3.1) and several large-deviation bounds for weakly dependent multinomial data (Lemmas 6.3-6.6).

Proof of Theorem 2.1: Write $Z = D - D_0$. First, we introduce a quantity $\Delta(Z, D_0)$ to capture the "noise" level and derive an upper bound for it. For $1 \leq j \leq p$, recall that a_j' denotes the j-th row of A and $h_j = \|a_j\|_1$; moreover, denote by $\hat{\Xi}_j'$ and Ξ_j' the j-th row of $\hat{\Xi} = [\hat{\xi}_1, \ldots, \hat{\xi}_K]$ and $\Xi = [\xi_1, \ldots, \xi_K]$, respectively. Noting that the eigenvectors are determined up to a multiple of ± 1, we let $\mathcal{O}_K$ be the set of all diagonal matrices $\Omega = \mathrm{diag}(\omega_1, \ldots, \omega_K)$ such that $\omega_k \in \{\pm 1\}$. Define

$$\Delta(Z, D_0) \equiv \max\{\Delta_1(Z, D_0), \ \ \Delta_2(Z, D_0)\}, \tag{11}$$

where

$$\Delta_1(Z, D_0) = \min_{\Omega \in \mathcal{O}_K} \max_{1 \leq j \leq p} \big\{h_j^{-1/2} \|\Omega\hat{\Xi}_j - \Xi_j\|\big\},$$
$$\Delta_2(Z, D_0) = \max_{1 \leq j \leq p} \big\{h_j^{-1} |M(j,j) - M_0(j,j)|\big\}.$$

We then bound $\Delta(Z, D_0)$. The analysis of $\Delta_2(Z, D_0)$ is a standard application of the Bernstein's inequality, and we relegate it to Section 6.2. To study $\Delta_1(Z, D_0)$, we need an element-wise bound for the empirical singular vectors $\hat{\xi}_1, \ldots, \hat{\xi}_K$. The following theorem is proved in Section 3.1.

THEOREM 3.1 (Perturbation of singular vectors). *Consider models (4)-(5), where (6)-(7) hold. Suppose* $\min\{N, p\} \geq \log(n)$ *and* $Nh_{\min} \geq \log^2(n)$. *With probability* $1 - o(n^{-3})$, *there exists* $\Omega \in \mathcal{O}_K$ *such that*

$$\|\Omega\hat{\Xi}_j - \Xi_j\| \leq C\Big(1 + \frac{p}{h_{\min}N^{3/2}}\Big)\sqrt{\frac{\log(n)}{Nn}}\big(1 + \sqrt{ph_j}\big), \quad \textit{for all } 1 \leq j \leq p.$$

If, furthermore, $h_{\min} \geq C^{-1}/p$ *and* $p \leq CN^{3/4}$ *for a constant* $C > 0$, *then with probability* $1 - o(n^{-3})$, *there exists* $\Omega \in \mathcal{O}_K$ *such that*

$$\frac{\|\Omega\hat{\Xi}_j - \Xi_j\|}{\sqrt{h_j}} \leq C\sqrt{\frac{p\log(n)}{Nn}}, \qquad \textit{for all } 1 \leq j \leq p.$$

By Theorem 3.1 and Lemma 6.3, we find that with probability $1 - o(n^{-3})$,

$$\Delta(Z, D_0) \leq C(Nn)^{-1/2}\sqrt{p\log(n)}. \tag{12}$$

Next, we characterize how the estimation error depends on the quantity $\Delta(Z, D_0)$. The following theorem is a non-stochastic result, and it is proved in Section 3.2.

THEOREM 3.2 (Non-stochastic error bounds). *Consider models (4)-(5), where (6)-(9) hold and* $\Delta(Z, D_0)$ *is defined as in (11). Suppose the tuning parameters in our method are* $t = \infty$ *and* $L \geq L_0 + K$, *where* L_0 *is as in (9). There are constants* $C, c > 0$ *such that, if* $\Delta(Z, D_0) \leq c$, *then there exists a permutation matrix* $T \in \mathcal{P}_K$ *such that*

$$\|T\hat{a}_j - a_j\|_1 \leq \|a_j\|_1 \cdot C\Delta(Z, D_0), \qquad \textit{for all } 1 \leq j \leq p.$$

The first claim of Theorem 2.1 follows from (12) and Theorem 3.2. Using the definition of $\mathcal{L}(\hat{A}, A)$ and the fact that $\sum_{j=1}^{p}\|a_j\|_1 = K$, we also obtain the second claim. □

3.1. *Perturbation analysis of the singular vectors (proof of Theorem 3.1).* Our main tool is the following lemma which characterizes the perturbation of eigenvectors when a matrix is perturbed. It can be viewed as a generalization of the sin-theta theorem (Davis and Kahan, 1970): The sin-theta theorem yields a bound on $\|\hat{\xi}_k - \xi_k\|$, while Lemma 3.1 provides a bound for each coordinate of $(\hat{\xi}_k - \xi_k)$.

LEMMA 3.1. *Let G and G_0 be two $p \times p$ symmetric matrices, where the rank of G_0 is K. For $1 \le k \le K$, let (λ_k, u_k) be the k-th eigenpair*[7] *of G_0, and let $(\hat{\lambda}_k, \hat{u}_k)$ be the k-th eigenpair of G. Let $E = G - G_0$, and denote by E_j' the j-th row of E, for $1 \le j \le p$. Suppose for some $a \in (0,1)$,*

$$\min_{1\le k\le K} |\lambda_k| \ge a\|G_0\|, \quad \min_{1\le k\le K-1} |\lambda_k - \lambda_{k+1}| \ge a\|G_0\|, \quad \|E\| \le (a/2)\|G_0\|.$$

Then, there exist $\omega_1, \ldots, \omega_K \in \{\pm 1\}$ such that, for all $1 \le j \le p$,

$$\max_{1\le k\le K} |\omega_k \hat{u}_k(j) - u_k(j)| \le \frac{C(a,K)}{\|G_0\|}\left(\|E\| \sum_{\ell=1}^{K} |u_\ell(j)| + \|E_j\|\right),$$

where $C(a,K) > 0$ is a constant that only depends on a and K.

Remark. In the conclusion of Lemma 3.1, if we take the sum of squares of both sides for $j = 1, 2, \ldots, p$, it yields that $\|\omega_k \hat{u}_k - u_k\| \le C\|G_0\|^{-1}(\|E\| + \|E\|_F)$. The first term matches what is given by the sin-theta theorem, suggesting that this term is tight (however, the sin-theta theorem fails to "allocate" the error to individual coordinates). There is space for improving the second term, but for our purpose of proving Theorem 3.1, this term is good enough.

We now prove Theorem 3.1. In our settings, $\hat{\xi}_k$ is the k-th singular vector of $M^{-1/2}D$ and ξ_k is the k-th singular vector of $M_0^{-1/2}D_0$. Equivalently, we can view $\hat{\xi}_k$ and ξ_k as the respective k-th leading eigenvector of

$$G \equiv M^{-1/2}DD'M^{-1/2} - \frac{n}{N}I_p \quad \text{and} \quad G_0 \equiv (1 - \frac{1}{N})M_0^{-1/2}D_0D_0'M_0^{-1/2}.$$

First, we conduct eigen-analysis on G_0. The next two lemmas study the eigenvalues and eigenvectors, respectively.

LEMMA 3.2. *Suppose the conditions of Theorem 3.1 hold. Let $\lambda_1, \ldots, \lambda_K$ be the nonzero eigenvalues of G_0. There exists a constant $C > 1$ such that*

$$C^{-1}n \le \min_{1\le k\le K} |\lambda_k| \le \max_{1\le k\le K} |\lambda_k| \le Cn, \qquad \min_{1\le k\le K-1} |\lambda_k - \lambda_{k+1}| \ge C^{-1}n.$$

LEMMA 3.3. *Suppose the conditions of Theorem 3.1 hold. Let $\xi_1, \ldots, \xi_K$ be the K leading eigenvectors of G_0. There exists a constant $C > 0$ such that*

$$\sum_{\ell=1}^{K} |\xi_\ell(j)| \le C\sqrt{h_j}, \qquad \text{for all } 1 \le j \le p.$$

[7] I.e., λ_k is the k-th eigenvalue and u_k is the associated unit-norm eigenvector.

Next, we study the matrix $G - G_0$. The next two lemmas provide bounds on the spectral norm and the ℓ_2-norm of an individual column, respectively.

LEMMA 3.4. *Under the conditions of Theorem 3.1, with probability* $1 - o(n^{-3})$, *for all* $1 \le j \le p$,

$$\|(G - G_0)e_j\| \le C(1 + \sqrt{ph_j}) \cdot N^{-1/2}\sqrt{n\log(n)}.$$

LEMMA 3.5. *Under the conditions of Theorem 3.1, with probability* $1 - o(n^{-3})$,

$$\|G - G_0\| \le C\Big(1 + \frac{p}{h_{\min}N^{3/2}}\Big) \cdot N^{-1/2}\sqrt{np\log(n)}.$$

Remark. The key of proving Lemmas 3.4-3.5 is to study the noise matrix $Z = D - D_0$. This is a random matrix whose columns are generated from multinomial distributions. Standard Random Matrix Theory does not apply because the entries of Z are weakly dependent. In Section 6.2, we carefully investigate properties of Z and state several lemmas that are useful for the proof of Lemmas 3.4-3.5.

Proof of Theorem 3.1: The first claim follows from plugging Lemmas 3.2-3.5 into Lemma 3.1. To get the second claim, we note that when $h_{\min} \ge C^{-1}/p$ and $CN^{3/4} \ge p$, it holds that $\sqrt{ph_j} \ge C^{-1/2}$ and $h_{\min}N^{3/2} \ge C^{-3}p$; hence, the bound for $\|\Omega\hat{\Xi}_j - \Xi_j\|$ now becomes $C(Nn)^{-1/2}\sqrt{ph_j\log(n)}$. □

3.2. *Non-stochastic analysis of our method (proof of Theorem 3.2).* We first study the the vertices hunting algorithm in Section 1.2, which takes the input $\hat{r}_1, \ldots, \hat{r}_p$ and outputs the estimated vertices $\hat{v}_1^*, \ldots, \hat{v}_K^*$. Our study of this algorithm is based on a key observation that each $\hat{r}_j$ is within a distance of $C\Delta_1(Z, D_0)$ to the corresponding r_j (as justified below). Consequently, we can estimate the K vertices up to an ℓ_2-error of $C\Delta_1(Z, D_0)$.

LEMMA 3.6. *Suppose the conditions of Theorem 3.2 hold. Let* $\mathcal{O}_{K-1}$ *be the set of all* $(K-1)\times(K-1)$ *diagonal matrices whose diagonals are equal to* ± 1. *There exists* $\Omega^* \in \mathcal{O}_{K-1}$ *such that*

$$\|\Omega^*\hat{r}_j - r_j\| \le C\Delta_1(Z, D_0), \qquad \textit{for all } 1 \le j \le p.$$

Furthermore,

$$Err_{VH} \equiv \min_{\kappa:\ a\ permutation\ on\ \{1,\ldots,K\}} \Big\{\max_{1\le k\le K} \|\Omega^*\hat{v}_k^* - v_{\kappa(k)}^*\|\Big\} \le C\Delta_1(Z, D_0).$$

Remark. Our vertices hunting algorithm is similar to the vertices hunting algorithm in Jin, Ke and Luo (2016), but the analysis is very different. Their analysis (when translated to our settings) requires that each topic has a constant fraction of anchor words, which is not satisfied in many real data. The main reason that we can remove this restriction is due to the stronger result on empirical singular vectors (see Theorem 3.1); such result allows us to bound each individual $\|\hat{r}_j - r_j\|$, while the analysis in Jin, Ke and Luo (2016) is built upon a bound for $\sum_{j=1}^{p} \|\hat{r}_j - r_j\|^2$.

We then investigate how the estimation error of our method depends on the noise level $\Delta(Z, D_0)$ and the accuracy of vertices hunting Err_{VH}. Then, Theorem 3.2 immediately follows from Lemmas 3.6-3.7.

LEMMA 3.7. *Under the conditions of Theorem 3.2, there exists a permutation matrix $T \in \mathcal{P}_K$ such that*

$$\|T\hat{a}_j - a_j\|_1 \leq \|a_j\|_1 \cdot C\big[\Delta(Z, D_0) + Err_{VH}\big]$$

Remark. Lemma 3.7 is not tied to any specific vertices hunting algorithm, and it holds even when the assumptions (8)-(9) are not satisfied. This is why we keep Err_{VH} in the bound rather than plugging in $Err_{VH} \leq C\Delta_1(Z, D_0)$.

Below, we give an outline of the proof of Lemma 3.7. The rigorous proof can be found in Section 6.

Proof sketch: From now on, for notation simplicity, we omit the permutation $\kappa(\cdot)$ in all places. Introduce a matrix

$$\hat{Q} = \begin{pmatrix} 1 & \ldots & 1 \\ \hat{v}_1^* & \ldots & \hat{v}_K^* \end{pmatrix},$$

and let $Reg(\cdot)$ denote the operation of regularizing a K-dimensional vector to a weight vector (i.e., truncating the negative entries and renormalizing them to have a unit ℓ_1-norm). Given the estimated vertices, our method can be equivalently expressed as follows: Compute $\hat{A}^* = [\hat{a}_1^*, \ldots, \hat{a}_p^*]'$ by

$$\hat{a}_j^* = \sqrt{M(j,j)} \cdot \hat{\xi}_1(j) \cdot Reg\big(\hat{Q}^{-1}\hat{r}_j\big), \qquad 1 \leq j \leq p, \tag{13}$$

and obtain $\hat{A}$ by normalizing each column of $\hat{A}^*$ to have a unit ℓ_1-norm.

The key is to study the error in $\hat{a}_j^*$. It can be shown that

$$\hat{a}_j^* \approx \sqrt{M(j,j)} \cdot \hat{\xi}_1(j) \cdot \hat{Q}^{-1}\hat{r}_j.$$

Let Q be the matrix with each $\hat{v}_k^*$ replaced by v_k^* in the definition of $\hat{Q}$; so, Q is a population counterpart of $\hat{Q}$. We then define a population counterpart of $\hat{a}_j^*$ by

$$a_j^* = \sqrt{M_0(j,j)} \cdot \xi_1(j) \cdot Q^{-1} r_j.$$

To bound $\|\hat{a}_j^* - a_j^*\|_1$, we look at the four quantities: $M(j,j)$, $\hat{\xi}_1(j)$, $\hat{Q}$ and $\hat{r}_j$. Using the definition (11) and the facts that $\xi_1(j) \asymp \sqrt{h_j}$ and that $M_0(j,j) \asymp h_j$ (see Lemma 6.2 and equation (49)), we have

$$\frac{|\hat{\xi}_1(j) - \xi_1(j)|}{|\xi_1(j)|} \le C\Delta_1(Z, D_0), \qquad \frac{|M(j,j) - M_0(j,j)|}{M_0(j,j)} \le C\Delta_2(Z, D_0).$$

Moreover, by Lemma 3.6 and the way $\hat{Q}$ is defined, up to a sign change for the coordinates of $\hat{r}_j$,

$$\|\hat{r}_j - r_j\| \le C\Delta_1(Z, D_0), \qquad \|\hat{Q} - Q\| \le Err_{VH}.$$

Combining the above results, we can prove that (note: a permutation matrix T has been omitted)

$$\frac{\|\hat{a}_j^* - a_j^*\|_1}{\|a_j^*\|_1} \le C\big[\Delta_1(Z, D_0) + \Delta_2(Z, D_0) + Err_{VH}\big]. \tag{14}$$

Given (14), it remains to study the effect of conducting a column-wise normalization on $\hat{A}^*$; note that such a normalization "aggregates" the errors in $\hat{a}_1^*, \ldots, \hat{a}_p^*$. We deal with this carefully in the proof of Lemma 3.7.

4. Simulations. We study the numerical performance of our method, where Section 4.1 contains experiments on simulated data and Section 4.2 contains experiments on semi-synthetic data from the AP and NIPS corpora. We call our method Topic-SCORE (or T-SCORE).

In all experiments below, we assume the number of topics K is known. Our method has two tuning parameters (t, L). We set $t = \infty$ and $L = 10 \times K$. We compare our method with three different methods: LDA (Blei, Ng and Jordan, 2003), AWR (Arora et al., 2013), and TSVD (Bansal, Bhattacharyya and Kannan, 2014). We implement LDA using the R package *lda*, with the default Dirichlet priors ($\alpha = \beta = 0.1$). We implement AWR using the Python code downloaded from `http://people.csail.mit.edu/moitra/software.html`. We implement TSVD using the matlab code downloaded from `http://thetb.github.io/tsvd/`.

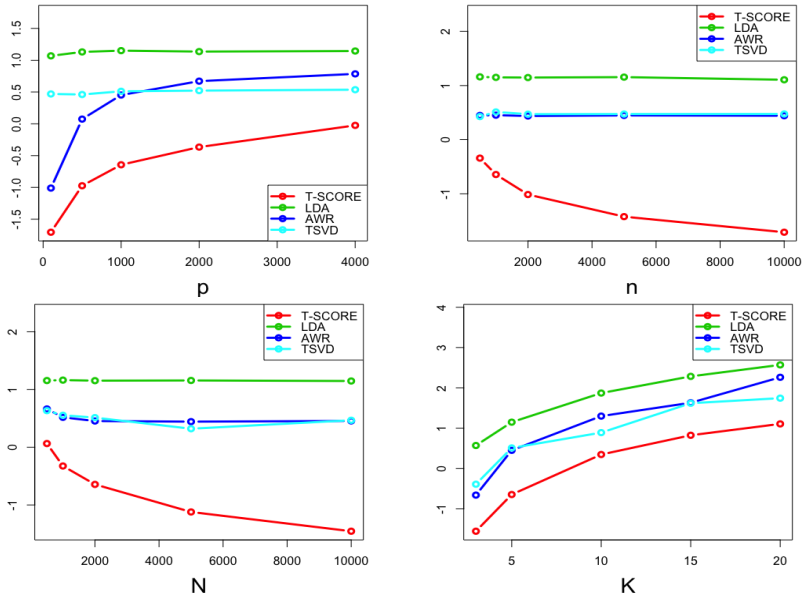

FIG 4. *Experiment 1. The y-axis is* $\log(\mathcal{L}(\hat{A}, A))$*, and* (p, n, N, K) *represent the vocabulary size, number of documents, document length, and number of topics, respectively.*

4.1. *Synthetic data.* Given parameters $\{p, n, N, K, m_p, \delta_p, m_n\}$, we generate the text corpus D as follows:

- Generate the topic matrix A: For $1 \le k \le K$, let each of the $[(k-1)m_p+1]$-th row to the (km_p)-th row equal to $\delta_p e_k'$, where $e_1, \ldots, e_K$ are the standard basis vectors of $\mathbb{R}^K$. For the remaining $(p-Km_p)$ rows, we first generate all entries *iid* from $Unif(0,1)$, and then normalize each column of the $(p-Km_p) \times K$ sub-matrix to have a sum of $(1-m_p\delta_p)$.
- Generate the document matrix W: For $1 \le k \le K$, let each of the $[(k-1)m_n+1]$-th column to the (km_n)-th column equal to e_k. For the remaining columns, we first generate all entries *iid* from $Unif(0,1)$, and then normalize each column to have a sum of 1.
- Generate the text corpus D using the model (5) and (4).

With this data generating process, there are m_p anchor words and m_n pure documents for each topic, and all the anchor words have a separability of δ_p. For each parameter setting, we independently generate 200 data sets and report the average $\mathcal{L}(\hat{A}, A)$ for all four methods.

Experiment 1: Various settings of (p, n, N, K). We fix a basic setting where $(p, n, N, K, m_p, \delta_p, m_n) = (1000, 1000, 2000, 5, p/100, 1/p, n/100)$. In the four

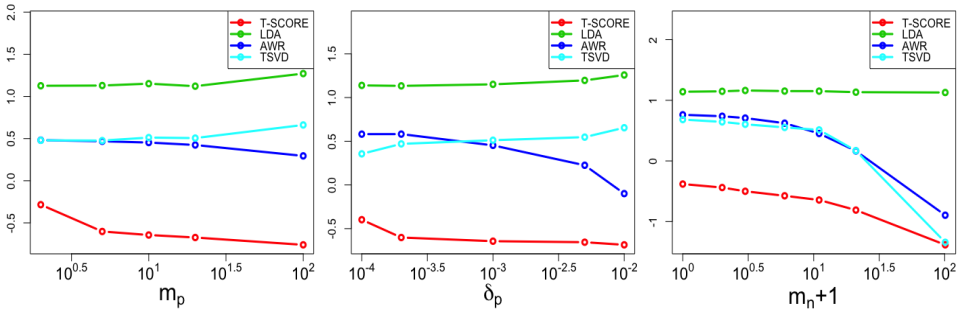

FIG 5. *Experiment 2. The y-axis is* $\log(\mathcal{L}(\hat{A}, A))$*, and* (m_p, δ_p, m_n) *represent the number of anchor words, separability of anchor words, and number of pure documents, respectively.*

sub-experiments, we vary one model parameter and keep the other parameters the same as in the basic setting. The results are shown in Figure 4. In all the settings, our method yields the smallest estimation error among all four methods. Furthermore, we have the following observations: (i) As n or N increases, our method is the only one whose estimation error exhibits a clear decreasing trend. It suggests that our method can take advantage of including *more* documents and having *longer* documents. (ii) As K increases, the estimation errors of all four methods increase, suggesting that the problem becomes more challenging for larger K. (iii) As p increases, the estimation errors of our method and AWR both increase, while the estimation errors of LDA and TSVD remain relatively stable; however, even for large p (e.g., $p = 4000$), still, our method significantly outperforms LDA and TSVD.

Experiment 2: Anchor words and pure documents. We fix the same basic setting as in Experiment 1 and vary one parameter of (m_p, δ_p, m_n) in each sub-experiment. The results are shown in Figure 5.

First, we look at the effect of anchor words. From the left panel of Figure 5, as m_p (number of anchor words per topic) increases, the estimation error of our method has considerably decreased, suggesting that our method can take advantage of having multiple anchor words. Even with $m_p = 2$, our method still outperforms the other methods. From the middle panel of Figure 5, as δ_p (separability of anchor words) increases, the estimation errors of AWR and our method both decrease, and they both outperform LDA and TSVD; with the same separability, our method always outperforms AWR. Furthermore, as long as δ_p is larger than 2×10^{-4}, our method is relatively insensitive to δ_p; this is consistent with the theory in Section 2.

Second, we look at the effect of pure documents. From the right panel of Figure 5, as m_n (number of pure documents) increases, the performance

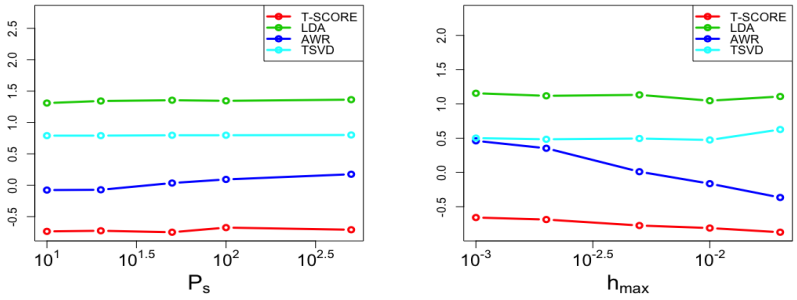

FIG 6. *Experiment 3. The y-axis is* $\log(\mathcal{L}(\hat{A}, A))$. *Left panel: the setting of Zipf's law. Right panel: the setting of two scales. The word heterogeneity increases as either* P_s *decreases or* $h_{\max}$ *increases.*

of all methods except LDA improves. The improvement on TSVD is especially significant; this is because TSVD relies on the existence of nearly-pure documents (which they called "dominant admixtures"). When $m_n < 100$, our method has a significant advantage over TSVD; when $m_n = 100$, the performance of our method is similar to that of TSVD.

Experiment 3: Heterogenous words. We study "heterogenous" settings where some words are much more frequent than the others. Fix $(p, n, N, K, m_p, \delta_p, m_n) = (1000, 1000, 2000, 5, p/100, 1/p, n/100)$. We generate the first Km_p rows of A in the same way as before and generate the remaining $(p - Km_p)$ rows using two different settings below:

- *Setting 1: Zipf's law.* Given $P_s > 0$, we first generate $A(j,k)$ from the exponential distribution with mean $(P_s + j)^{-1.07}$, independently for all $1 \le k \le K$, $Km_p + 1 \le j \le p$, and then normalize each column of the $(p - Km_p) \times K$ matrix to have a sum of $(1 - m_p\delta_p)$. Under this setting, the word frequencies of each topic roughly follow a Zipf's law with P_s stopping words. A smaller P_s corresponds to larger heterogeneity.
- *Setting 2: Two scales.* Given $h_{\max} \in [1/p, 1)$, first, we generate $\{A(j,k) : 1 \le k \le K, Km_p < j \le Km_p + n_{\max}\}$ *iid* from $Unif(0, h_{\max})$, where $n_{\max} = \lfloor (1 - m_p\delta_p)/(2h_{\max}) \rfloor$. Next, we define $n_{\min} = p - Km_p - n_{\max}$ and $h_{\min} = (1 - m_p\delta_p - h_{\max}n_{\max})/n_{\min}$ and generate $\{A(j,k) : 1 \le k \le K, Km_p + n_{\max} < j \le p\}$ *iid* from $Unif(0, h_{\min})$. Last, we normalize each column of the $(p - Km_p) \times K$ matrix to have a sum of $(1 - m_p\delta_p)$. Under this setting, the word frequencies of each topic are in two distinct scales, characterized by $h_{\max}$ and $h_{\min}$, respectively.

We then generate (W, D) in the same way as before. The results are shown in Figure 6. Our method always yields the smallest estimation errors. In-

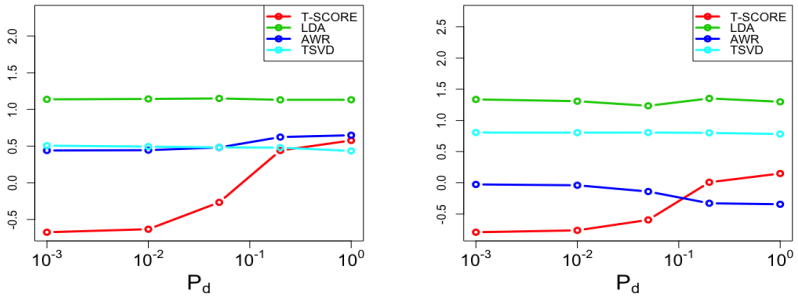

FIG 7. *Experiment 4. The y-axis is* $\log(\mathcal{L}(\hat{A}, A))$. *As* P_d *increases, the almost-anchor words are less anchor-like. Left panel: the homogeneous setting. Right panel: the heterogeneous setting.*

terestingly, in Setting 2, the performance of AWR improves with increased heterogeneity; see the right panel of Figure 6.

Experiment 4: No exact anchor words. Fix $(p, n, N, K, m_p, \delta_p, m_n, P_s) = (1000, 1000, 2000, 5, p/100, 1/p, n/100, p/20)$. We generate A using two different settings below:

- *Setting 1: Homogeneous words.* Given $P_d \in [0, 1]$, for $1 \le k \le K$, let each of the $[(k-1)m_p + 1]$-th row to the (km_p)-th row equal to $\delta_p \tilde{e}_k'$, where $\tilde{e}_k(j) = 1\{j = k\} + P_d 1\{j \ne k\}$, $1 \le j \le K$. For the remaining $(p - Km_p)$ rows, we first generate all entries *iid* from $Unif(0, 1)$, and then normalize each column of the $(p - Km_p) \times K$ sub-matrix to have a sum of $[1 - m_p\delta_p - m_p\delta_p(K-1)P_d]$.
- *Setting 2: Heterogenous words.* Given $P_d \in [0, 1]$, first, we generate $A(j, k)$ from the exponential distribution with mean $(P_s + j)^{-1.07}$, independently for all $1 \le k \le K$, $1 \le j \le p$; second, for each $1 \le k \le K$, we randomly select m_p rows from all the rows whose largest entry is the k-th entry, and for these selected rows, we keep the k-th entry and multiply the other entries by P_d; last, we renormalize each column of A to have a sum of 1.

We then generate (W, D) in the same way as before. In both settings, there are m_p almost-anchor words for each topic. Moreover, a smaller P_d means that the almost-anchor words are more similar to anchor words; in the special case of $P_d = 0$, they become exact anchor words.

The results are shown in Figure 7. In both settings, our method yields the smallest estimation errors in a wide range of P_d, suggesting that our method has reasonable performance even without exact anchor words. In Setting 1, when $P_d = 1$, TSVD yields the best performance and the performance of our

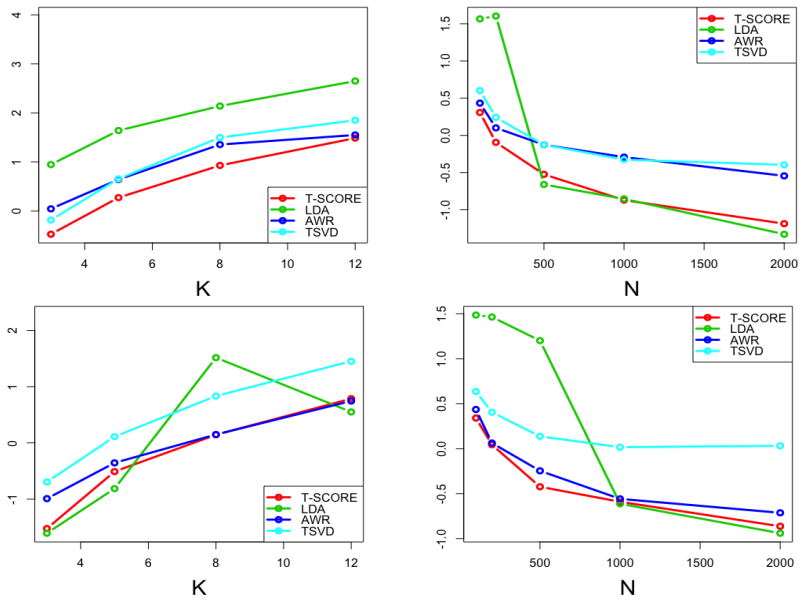

FIG 8. *Semi-synthetic experiments. The y-axis is* $\log(\mathcal{L}(\hat{A}, A))$. *Top panels: the AP corpus* ($n = 2135, p = 5188$). *Bottom panels: the NIPS corpus* ($n = 1417, p = 2508$).

TABLE 5
Computation time on the semi-synthetic data ($N = 2000, K = 5$).

Method	Software	AP data (in second)	NIPS data (in second)
Topic-SCORE	R	1.04	0.29
LDA	R	378.04	395.14
AWR	Python	112.62	36.68
TSVD	MATLAB	4.41	1.61

method is slightly worse than that of TSVD. In Setting 2, when $P_d > 0.1$, our method is better than LDA and TSVD but is worse than AWR. Interestingly, although AWR relies on the existence of anchor-like words, its performance actually improves as P_d increases; the reason is unclear to us.

4.2. *Semi-synthetic data from the AP and NIPS corpora.* Semi-synthetic experiments are commonly used in the literature of topic model estimation. Given a real data set with n documents written on a vocabulary of p words, with pre-specified $(K, N_1, \ldots, N_n)$, we first run LDA by assuming K topics; next, using the posterior of (A, W) obtained from LDA, we generate n new documents such that document i has N_i words, $1 \leq i \leq n$. We took the AP data set (Harman, 1993) and the NIPS data set (Perrone et al., 2016)

and preprocessed them by removing stop words and keeping the 50% most frequent words and 95% longest documents. For each data set, we conducted two experiments: In the first experiment, $(N_1, \ldots, N_n)$ are the same as in the original data set and K varies in $\{3, 5, 8, 12\}$. In the second experiment, $K = 5$ and $N_i = N$ for all $1 \leq i \leq n$, with N varying in $\{100, 200, 500, 1000, 2000\}$.

The results are shown in Figure 8. Our method outperforms TSVD and AWR in almost all settings and outperforms LDA in many settings (note that the data generating process favors LDA). In Table 5, we compare the computing time of different methods. Our method is much faster than LDA and AWR and is comparable with TSVD.

5. Discussion. We propose a new SVD approach to topic estimation, at the heart of which is the discovery of a low-dimensional simplex structure associated with the entry-wise eigen-ratios. The method is successfully applied to real applications, and is shown to be asymptotically optimal.

We recognize both that SVD is a powerful tool for dimension reduction and noise reduction and that SVD faces challenges in many modern applications. Our approach adapts SVD for modern uses, and for this reason, it is connected to many recent works on a high level. These include but are not limited to the works on sparse PCA (Berthet and Rigollet, 2013; Wang, Berthet and Samworth, 2016; Han and Liu, 2014; Vu and Lei, 2013; Arias-Castro, Lerman and Zhang, 2017), the works on IF-PCA (Jin and Wang, 2016), the works on factor models (Fan, Fan and Lv, 2008; Fan, Liao and Mincheva, 2011), and the works on SCORE (Jin, 2015; Ji and Jin, 2016).

Our method has two carefully-designed normalizations: a pre-SVD normalization by the matrix $M^{-1/2}$ and a post-SVD normalization by taking entry-wise eigen-ratios. The former improves the performance of SVD, and is crucial for the minimax optimality. The latter is crucial for the simplex construction; it is connected to Jin (2015) on social networks, on a high level, but two papers deal with very different problems, where the models and analysis are very different.

The method is convenient to use and its computation is reasonably fast, even for large data sets. For this reason it can be used as a starting point for methods that are both more complicated and computationally slower. These include but are not limited to the popular Bayesian approaches (Airoldi et al., 2008; Blei and Lafferty, 2007).

The core idea is very flexible and can be extended to many different settings, such as Nonnegative Matrix Factorization (NMF) (Paatero and Tapper, 1994; Lee and Seung, 1999; Donoho and Stodden, 2004). NMF is useful for many scientific projects, and it aims to factorize a given matrix to

be approximately equal to the product of two low-rank nonnegative matrices. In a manuscript Ke and Wang (2017), we extend the main idea of this paper and develop a new SVD-based algorithm for NMF.

Our method and theory can be modified to accommodate more general settings. When the topic vectors are sparse, we only need to modify the SVD part in our method, say, by conducting a pre-screening on words or replacing it by sparse PCA methods (Zou, Hastie and Tibshirani, 2006). In this paper, we assume the number of topics K is fixed and known. How to estimate K is a challenging problem (Owen and Wang, 2016; Saldana, Yu and Feng, 2017). Also, it is possible to extend our method and theory to the case where K grows to infinity (the minimax rate will then depend on K). We leave this to future work.

6. Proofs.

6.1. *Preliminary I: The two matrices of entry-wise ratios.* First, we consider the matrix $V^* \in \mathbb{R}^{K,K-1}$. It is obtained from taking the entry-wise ratios of the matrix V, where V is defined by $\Xi = AV$ (if it exists). Write $V = [V_1, \ldots, V_K]$ and $V^* = [v_1^*, \ldots, v_K^*]'$.

LEMMA 6.1. *Under the assumptions of Theorem 2.1, the following statements are true:*

- *Fixing the choice of Ξ, there is a unique non-singular matrix $V \in \mathbb{R}^{K,K}$ such that $\Xi = M_0^{-1/2}AV$; moreover, $(VV')^{-1} = A'M_0^{-1}A$.*
- *All the entries of V_1 have the same sign; moreover, $C_1^{-1} \le |V_1(k)| \le C_1$ for all $1 \le k \le K$.*
- *$\mathcal{S}_K^* = \mathcal{S}(v_1^*, \ldots, v_K^*)$ is a non-degenerate simplex; moreover, the volume of $\mathcal{S}_K^*$ is lower bounded by C_2^{-1} and upper bounded by C_2.*
- $\max_{1\le k\le K} \|v_k^*\| \le C_3$.
- *$C_4^{-1} \le \|v_k^* - v_\ell^*\| \le C_4$ for all $1 \le k \ne \ell \le K$.*

Here, C_1-C_4 are positive constants satisfying that $C_1, C_2, C_4 > 1$.

Next, we consider the matrix R. It is obtained from taking the entry-wise ratios of the matrix $\Xi = [\xi_1, \ldots, \xi_K]$. For $1 \le j \le p$, recall that a_j' denotes the j-th row of A, and $\tilde{a}_j = h_j^{-1}a_j$, where $h_j = \|a_j\|_1$. Write $R = [r_1, \ldots, r_p]'$.

LEMMA 6.2. *Under the assumptions of Theorem 2.1, the following statements are true:*

- *We can choose the sign of ξ_1 such that all the entries are positive and that $C_5^{-1}\sqrt{h_j} \le \xi_1(j) \le C_5\sqrt{h_j}$ for all $1 \le j \le p$.*

- $\max_{1\le j\le p}\|r_j\| \le C_6$.
- $C_7^{-1}\|\tilde{a}_i - \tilde{a}_j\| \le \|r_i - r_j\| \le C_7\|\tilde{a}_i - \tilde{a}_j\|$, *for all* $1 \le i, j \le p$.

Here, C_5-C_7 *are positive constants satisfying that* $C_5, C_7 > 1$.

Lemmas 6.1-6.2 are proved in Section 7.

6.2. *Preliminary II: The noise matrix* $Z = D - D_0$. The distribution of Z is characterized by the model (4). Let $\{h_j\}_{j=1}^p$ be as defined in Section 2.1, and write $h_{\max} = \max_{1\le j\le p} h_j$ and $h_{\min} = \max_{1\le j\le p} h_j$. Write

$$Z = [z_1, z_2, \ldots, z_n] = [Z_1, Z_2, \ldots, Z_p]'.$$

The following lemma is about the diagonal matrix $M - M_0 = n^{-1}\mathrm{diag}(Z1_n)$.

LEMMA 6.3. *Suppose* $Nnh_{\min}/\log(n) \to \infty$. *With probability* $1 - o(n^{-3})$,

$$|M(j,j) - M_0(j,j)| \le C(Nn)^{-1/2}\sqrt{h_j\log(n)}, \qquad \text{for all } 1 \le j \le p.$$

The following lemma is about the p-dimensional vector $M_0^{-1/2}ZW_k$, where W_k' denotes the k-th row of W, for $1 \le k \le K$.

LEMMA 6.4. *Suppose* $Nnh_{\min}/\log(n) \to \infty$. *With probability* $1 - o(n^{-3})$, *for all* $1 \le k \le K$,

$$|Z_j'W_k| \le CN^{-1/2}\sqrt{nh_j\log(n)}, \qquad \text{for all } 1 \le j \le p,$$
$$\|M_0^{-1/2}ZW_k\| \le CN^{-1/2}\sqrt{np\log(n)}.$$

The next two lemmas are about the $p \times p$ matrix ZZ', where Lemma 6.5 considers individual entries of it, and Lemma 6.6 studies its spectral norm.

LEMMA 6.5. *Suppose* $Nh_{\min}/\log(n) \to \infty$. *With probability* $1 - o(n^{-3})$,

$$|Z_j'Z_\ell - E[Z_j'Z_\ell]| \le CN^{-1}\sqrt{nh_jh_\ell\log(n)}, \qquad \text{for all } 1 \le j, \ell \le p.$$

LEMMA 6.6. *Suppose* $\min\{Nh_{\min}, p\}/\log(n+N) \to \infty$ *and* $p = O(n)$. *With probability* $1 - o(n^{-3})$,

$$\|M_0^{-1/2}(ZZ' - E[ZZ'])M_0^{-1/2}\| \le C\Big(\frac{1}{N} + \frac{p}{N^2h_{\min}}\Big)\sqrt{np}.$$

Lemmas 6.3-6.6 are proved in Section 7.

6.3. *Proof of Lemma 1.1.* Recall that V is the non-singular matrix such that $\Xi = M_0^{-1/2}AV$, where the existence and uniqueness of V are justified in Lemma 6.1. Moreover, by Lemmas 6.1-6.2, both V^* and R are well-defined; by their definitions, $V = \mathrm{diag}(V_1) \cdot [1_K, V^*]$ and $\Xi = \mathrm{diag}(\xi_1) \cdot [1_p, R]$. Combining the above, we have

$$\underbrace{\mathrm{diag}(\xi_1) \cdot [1_p, R]}_{\Xi} = M_0^{-1/2} A \cdot \underbrace{\mathrm{diag}(V_1) \cdot [1_K, V^*]}_{V}.$$

Equivalently,

$$[1_p, R] = \underbrace{[\mathrm{diag}(\xi_1)]^{-1} M_0^{-1/2} A \cdot \mathrm{diag}(V_1)}_{\Pi} \cdot [1_K, V^*]. \tag{15}$$

First, we show that each row of Π is indeed a weight vector. By Lemma 6.2, we can choose the sign of ξ_1 such that all its entries are positive; additionally, since $\xi_1 = AV_1$ and that each topic has a few anchor words, we find that the K entries of V_1 are also positive. Combining the above, Π is a non-negative matrix. Furthermore, it follows from (15) that $1_p = \Pi \cdot 1_K$, i.e., the row sums of Π are all equal to 1. Therefore, each row of Π is a weight vector. Second, using (15) again, $R = \Pi \cdot V^*$, which implies that each row of R is a convex combination of the rows of V^* with the weights being the corresponding row of Π. □

6.4. *Proof of Theorem 2.2.* We need a useful lemma:

LEMMA 6.7 (Kullback-Leibler divergence). *Let $D_0, \tilde{D}_0$ be two $p \times n$ matrices such that each column of them is a weight vector. Under Model (4), let $\mathbb{P}$ and $\tilde{\mathbb{P}}$ be the probability measures associated with D_0 and $\tilde{D}_0$, respectively, and let $KL(\tilde{\mathbb{P}}, \mathbb{P})$ be the Kullback-Leibler divergence between them. Suppose D_0 is a positive matrix. Let $\delta = \max_{1\le j\le p, 1\le i\le n} \frac{|\tilde{D}_0(j,i) - D_0(j,i)|}{D_0(j,i)}$ and assume $\delta < 1$. There exists a universal constant $C > 0$ such that*

$$KL(\tilde{\mathbb{P}}, \mathbb{P}) \le (1 + C\delta) N \sum_{i=1}^{n} \sum_{j=1}^{p} \frac{|\tilde{D}_0(j,i) - D_0(j,i)|^2}{D_0(j,i)}.$$

Below, we show Lemma 6.7. Write for short $a_{ji} = D_0(j,i)$, $\tilde{a}_{ji} = \tilde{D}_0(j,i)$, and $\delta_{ji} = \frac{\tilde{a}_{ji} - a_{ji}}{a_{ji}}$. Then, $\delta = \max_{i,j} |\delta_{ji}|$. Note that the KL-divergence between Multinomial(N, η_1) and Multinomial(N, η_2) is equal to $N \sum_{j=1}^{p} \eta_{1j} \log(\eta_{1j}/\eta_{2j})$. It follows that

$$KL(\tilde{\mathbb{P}}, \mathbb{P}) = N \sum_{i=1}^{n} \sum_{j=1}^{p} \tilde{a}_{ji} \log(1 + \delta_{ji}).$$

By Taylor expansion, $\log(1+\delta_{ji}) \le \delta_{ji} - \frac{1}{2}\delta_{ji}^2 + C\delta_{ji}^3$ for a constant $C > 0$. Moreover, since each column of D_0 and $\tilde{D}_0$ has a sum of 1, we have $\sum_{i,j} a_{ji} = \sum_{i,j} \tilde{a}_{ji}$, which implies that $\sum_{i,j} a_{ji}\delta_{ji} = 0$. As a result,

$$\begin{aligned} KL(\tilde{\mathbb{P}}, \mathbb{P}) &\le N \sum_{i,j} (a_{ji} + a_{ji}\delta_{ji})(\delta_{ji} - \frac{1}{2}\delta_{ji}^2 + C\delta_{ji}^3) \\ &= N \sum_{i,j} a_{ji}\delta_{ji} + N \sum_{i,j} a_{ji}\delta_{ji}^2 - \frac{N}{2} \sum_{i,j} a_{ji}\delta_{ji}^2 + O\Big(N \sum_{i,j} a_{ij}\delta_{ji}^3\Big) \\ &= \frac{N}{2} \sum_{i,j} a_{ji}\delta_{ji}^2 + O\Big(\delta \cdot N \sum_{i,j} a_{ij}\delta_{ji}^2\Big). \end{aligned}$$

Then, Lemma 6.7 follows.

We now show the claim. We first prove that the loss function $\mathcal{L}(\cdot,\cdot)$ is a semi-norm on $\Phi_{n,N,p}(K,c)$. It suffices to check the triangle inequality. We introduce some notations. For any A and $\tilde{A}$, let $|A - \tilde{A}|_1 = \sum_{k=1}^K \|A_k - \tilde{A}_k\|_1$. For any A and a permutation κ on $\{1,\ldots,K\}$, write $A^\kappa = [A_{\kappa(1)}, \ldots, A_{\kappa(K)}]$. Then, for any topic matrices $(A, \tilde{A}, \check{A})$, it is seen that $\mathcal{L}(A,\tilde{A}) = \min_\kappa |A^\kappa - \tilde{A}|_1 = \min_{\kappa_1,\kappa_2} |A^{\kappa_1} - \tilde{A}^{\kappa_2}|_1 \le \min_{\kappa_1,\kappa_2}(|A^{\kappa_1} - \check{A}|_1 + |\check{A} - \tilde{A}^{\kappa_2}|_1) = \min_{\kappa_1} |A^{\kappa_1} - \check{A}|_1 + \min_{\kappa_2} |\check{A} - \tilde{A}^{\kappa_2}| = \mathcal{L}(A,\check{A}) + \mathcal{L}(\check{A},\tilde{A})$.

Next, we apply Theorem 2.5 of Tsybakov (2009): If there exist $(A^{(0)}, W^{(0)})$, $(A^{(1)}, W^{(1)})$, …, $(A^{(J)}, W^{(J)}) \in \Phi_{n,N,p}(K,c)$ such that:

(i) $\mathcal{L}(A^{(j)}, A^{(k)}) \ge 2C_0\sqrt{\frac{p}{Nn}}$ for all $0 \le j \ne k \le J$,
(ii) $KL(\mathcal{P}_j, \mathcal{P}_0) \le \beta \log(J)$ for all $1 \le j \le J$,

where $C_0 > 0$, $\beta \in (0, 1/8)$, and $\mathcal{P}_j$ denotes the probability measure associated with $(A^{(j)}, W^{(j)})$, then

$$\inf_{\hat{A}} \sup_{(A,W)\in\Phi_{n,N,p}(K,c)} \mathbb{P}\Big(\mathcal{L}(\hat{A}, A) \ge C_0\sqrt{\tfrac{p}{Nn}}\Big) \ge \tfrac{\sqrt{J}}{1+\sqrt{J}}\Big(1 - 2\beta - \sqrt{\tfrac{2\beta}{\log(J)}}\Big).$$

As long as $J \to \infty$ as $(n,N,p) \to \infty$, the right hand side is lower bounded by a constant, and the claim follows.

What remains is to construct $(A^{(0)}, W^{(0)}), (A^{(1)}, W^{(1)}), \ldots, (A^{(J)}, W^{(J)})$ that satisfy (i) and (ii). Let $m = p/2$ if p is even and $m = (p-1)/2$ if p is odd. The Varshamov-Gilbert bound for the packing numbers (Tsybakov, 2009, Lemma 2.9) guarantees that there exist $J \ge 2^{m/8}$ and $\omega^{(0)}, \omega^{(1)}, \ldots, \omega^{(J)} \in \{0,1\}^m$ such that $\omega^{(0)} = (0,\ldots,0)$ and

$$\sum_{j=1}^m 1\{\omega_j^{(s)} \ne \omega_j^{(\ell)}\} \ge \frac{m}{8}, \qquad \text{for any } 0 \le s \ne \ell \le J.$$

Let $\alpha_n = \frac{16C_0}{K}\frac{1}{\sqrt{Nnp}}$. We construct $A^{(0)}, A^{(1)}, \ldots, A^{(J)}$ as follows:

$$A_k^{(s)} = \frac{1}{p}1_p + \alpha_n \begin{cases} (\omega^{(s)}, -\omega^{(s)})', & \text{if } p \text{ is even,} \\ (\omega^{(s)}, -\omega^{(s)}, 0)', & \text{if } p \text{ is odd,} \end{cases} \quad 1 \le k \le K, 0 \le s \le J.$$

We then choose $W^{(0)}, W^{(1)}, \ldots, W^{(J)}$ arbitrarily from $\mathcal{W}_p(K)$.

We check that both (i)-(ii) are satisfied. For any $0 \le s \ne \ell \le J$,

$$\mathcal{L}(A^{(s)}, A^{(\ell)}) = 2K\alpha_n\|\omega^{(s)} - \omega^{(\ell)}\|_1 \ge \frac{1}{4}K\alpha_n m \ge 2C_0\sqrt{\tfrac{p-1}{Nn}}.$$

So (i) is satisfied. We then verify (ii). Write $D_0^{(s)} = A^{(s)}W^{(s)}$. From the way $A^{(s)}$ is constructed, all the columns of $D_0^{(s)}$ are equal to $A_1^{(s)}$. As a result, $D_0^{(0)}$ is a matrix whose entries are all equal to $(1/p)$, $\frac{|D_0^{(s)}(j,i) - D_0^{(0)}(j,i)|}{D_0^{(0)}(j,i)} \le p\alpha_n \ll 1$ for $s \ne 0$, and

$$\sum_{i=1}^n\sum_{j=1}^p \frac{|D_0^{(s)}(j,i) - D_0^{(0)}(j,i)|^2}{D_0(j,i)} = np\alpha_n^2 \cdot 2K\|\omega^{(s)} - \omega^{(0)}\|_1 \le \frac{16^2C_0^2}{K}\frac{p}{N},$$

where we have used $\|\omega^{(s)} - \omega^{(0)}\|_1 \le m \le p/2$. It then follows from Lemma 6.7 that $KL(\mathcal{P}_j, \mathcal{P}_0) \le [1 + o(1)]\frac{16^2C_0^2}{K}p$. At the same time, $\log(J) \ge \frac{m}{8}\log(2) \ge \frac{\log(2)}{16}(p-1)$. By choosing C_0 appropriately small, (ii) is satisfied. The proof is now complete. □

6.5. *Proof of Lemma 3.1.* Fix $1 \le k \le K$. Without loss of generality, we assume $u_k'\hat{u}_k \ge 0$ and consider $\omega_k = 1$; if $u_k'\hat{u}_k < 0$, we consider $\omega_k = -1$, and the proof is similar. Introduce

$$\tilde{u}_k = \hat{\lambda}_k^{-1}\sum_{\ell=1}^K \lambda_\ell(u_\ell'\hat{u}_k)u_\ell, \qquad Q = (I_p - \hat{\lambda}_k^{-1}E)^{-1} - I_p.$$

By definition, $(G_0 + E)\hat{u}_k = \hat{\lambda}_k\hat{u}_k$, where $G_0 = \sum_{\ell=1}^K \lambda_\ell u_\ell u_\ell'$. It follows that

$$\sum_{\ell=1}^K \lambda_\ell(u_\ell'\hat{u}_k)u_\ell + E\hat{u}_k = \hat{\lambda}_k\hat{u}_k \tag{16}$$

As a result,

$$\hat{u}_k = (\hat{\lambda}_k I_p - E)^{-1}\sum_{\ell=1}^K \lambda_\ell(u_\ell'\hat{u}_k)u_\ell = (I_p + Q)\tilde{u}_k. \tag{17}$$

First, we bound $|\tilde{u}_k(j) - u_k(j)|$. It is seen that

$$|\tilde{u}_k(j) - u_k(j)| \le \hat{\lambda}_k^{-1} \sum_{\ell \neq k} |\lambda_\ell| |u_\ell' \hat{u}_k| \cdot |u_\ell(j)| + |\frac{\lambda_k}{\hat{\lambda}_k}(u_k'\hat{u}_k) - 1| \cdot |u_k(j)|$$

$$\equiv (I) + (II). \tag{18}$$

Consider (I). For $\ell \neq k$, if we multiply both sides of (16) by u_ℓ' from the left, using the mutual orthogonality of $u_1, \ldots, u_K$, we find that

$$\lambda_\ell(u_\ell'\hat{u}_k) + u_\ell' E\hat{u}_k = \hat{\lambda}_k(u_\ell'\hat{u}_k) \quad \Longrightarrow \quad u_\ell'\hat{u}_k = \frac{u_\ell' E\hat{u}_k}{\hat{\lambda}_k - \lambda_\ell}.$$

The numerator satisfies that $|u_\ell' E\hat{u}_k| \le \|E\|$. The denominator satisfies that $|\hat{\lambda}_k - \lambda_\ell| \ge |\lambda_k - \lambda_\ell| - |\hat{\lambda}_k - \lambda_k|$, where $|\lambda_k - \lambda_\ell| \ge a\|G_0\|$ and the Weyl's inequality yields $|\hat{\lambda}_k - \lambda_k| \le \|E\| \le (a/2)\|G_0\|$; hence, $|\hat{\lambda}_k - \lambda_\ell| \ge (a/2)\|G_0\|$. Combining the above gives

$$|u_\ell'\hat{u}_k| \le \frac{2}{a}\frac{\|E\|}{\|G_0\|}, \qquad \text{for } \ell \neq k. \tag{19}$$

Moreover, since $|\hat{\lambda}_k| \ge |\lambda_k| - \|E\| \ge (a/2)\|G_0\|$ and $|\lambda_\ell| \le \|G_0\|$, we have

$$(I) \le \frac{4\|E\|}{a^2\|G_0\|} \sum_{\ell \neq k} |u_\ell(j)|. \tag{20}$$

Consider (II). If we multiply both sides of (16) by $\hat{u}_k'$ from the left, then

$$\lambda_k(u_k'\hat{u}_k)^2 + \sum_{\ell \neq k} \lambda_\ell(u_\ell'\hat{u}_k)^2 + \hat{u}_k' E\hat{u}_k = \hat{\lambda}_k,$$

which implies

$$1 - (u_k'\hat{u}_k)^2 = \frac{\sum_{\ell \neq k} \lambda_\ell(u_\ell'\hat{u}_k)^2 + \hat{u}_k' E\hat{u}_k - (\hat{\lambda}_k - \lambda_k)}{\lambda_k}.$$

Here $|\hat{u}_k' E\hat{u}_k| \le \|E\|$ and $|\hat{\lambda}_k - \lambda_k| \le \|E\|$; moreover, by (19), $\sum_{\ell \neq k} \lambda_\ell(u_\ell'\hat{u}_k)^2 \le C\|E\|^2/\|G_0\| \le C\|E\|$. So the above suggests that $1-(u_k'\hat{u}_k)^2 \le C\|E\|/\|G_0\|$. Since $1 - u_k'\hat{u}_k \le 1 - (u_k'\hat{u}_k)^2$ (note that $u_k'\hat{u}_k \ge 0$), we immediately have

$$1 - u_k'\hat{u}_k \le C\frac{\|E\|}{\|G_0\|}. \tag{21}$$

It follows that

$$(II) \le \Big(\frac{|\hat\lambda_k - \lambda_k|}{\hat\lambda_k}|u_k'\hat u_k| + |1 - u_k'\hat u_k|\Big)\cdot|u_k(j)| \le \frac{C\|E\|}{\|G_0\|}|u_k(j)|. \tag{22}$$

Plugging (20) and (22) into (18) gives

$$|\tilde u_k(j) - u_k(j)| \le \frac{C\|E\|}{\|G_0\|}\sum_{\ell=1}^K |u_\ell(j)|. \tag{23}$$

Next, we bound $|\hat u_k(j) - \tilde u_k(j)|$. By (17), $\hat u_k - \tilde u_k = Q\tilde u_k$. Let q_j' be the j-th row of Q. It follows that

$$|\hat u_k(j) - \tilde u_k(j)| = |q_j'\tilde u_k| \le \|q_j\|\cdot\|\tilde u_k\| \le C\|q_j\|; \tag{24}$$

here the last inequality comes from (23): By the Cauchy-Schwarz inequality, $|\tilde u_k(j) - u_k(j)|^2 \le (C\|E\|/\|G_0\|)^2\cdot K\sum_{\ell=1}^K u_\ell^2(j)$; it follows that $\|\tilde u_k - u_k\|^2 \le C\|E\|^2/\|G_0\|^2$; hence, $\|\tilde u_k\|$ is bounded. In light of (24), we now bound $\|q_j\|$. By definition of Q, it satisfies the equality

$$Q = \hat\lambda_k^{-1}E + \hat\lambda_k^{-1}QE \quad\Longrightarrow\quad q_j' = \hat\lambda_k^{-1}E_j' + \hat\lambda_k^{-1}q_j'E.$$

It follows that

$$\|q_j\| \le \hat\lambda_k^{-1}\|E_j\| + (\|E\|/\hat\lambda_k)\|q_j\|.$$

As a result,

$$\|q_j\| \le \frac{\hat\lambda_k^{-1}\|E_j\|}{1 - (\|E\|/\hat\lambda_k)} \le C\hat\lambda_k^{-1}\|E_j\| \le \frac{C\|E_j\|}{\|G_0\|}. \tag{25}$$

Plugging (25) into (24) gives

$$|\hat u_k(j) - \tilde u_k(j)| \le \frac{C\|E_j\|}{\|G_0\|}. \tag{26}$$

The claim then follows from (23) and (26). □

6.6. *Proof of Lemmas 3.2-3.3.* First, consider Lemma 3.2. We use the linear algebra result: For any two matrices A and B, the nonzero eigenvalues of AB are the same as the nonzero eigenvalues of BA. Then, the nonzero eigenvalues of $G_0 = (1 - \frac{1}{N})M_0^{-1/2}D_0D_0'M_0^{-1/2} = (1 - \frac{1}{N})M_0^{-1/2}AWW'A'M_0^{-1/2}$ are the same as the nonzero eigenvalues of

$$(1 - \frac{1}{N})WW'A'M_0^{-1}A = (1 - \frac{1}{N})n\Sigma_W\widetilde\Sigma_A.$$

So $Eiggap(G_0) \geq (1-\frac{1}{N})Eiggap(\Sigma_W\tilde{\Sigma}_A) \geq (1-\frac{1}{N})c_2 n$, where we have used the assumption (7). This gives the second claim.

By (49), $c_1 h_j \leq M_0(j,j) \leq h_j$, for all $1 \leq j \leq p$. So, $1 \leq \lambda_{\min}(M_0^{-1/2}H^{1/2}) \leq \lambda_{\max}(M_0^{-1/2}H^{1/2}) \leq 1/\sqrt{c_1}$. It follows that ($s_{\min}(\cdot)$: the minimum singular value)

$$\begin{aligned} s_{\min}(G_0) &= (1-\tfrac{1}{N})\cdot s_{\min}\big(M_0^{-1/2}AWW'A'M_0^{-1/2}\big) \\ &\geq (1-\tfrac{1}{N})\cdot s_{\min}\big(H^{-1/2}AWW'A'H^{-1/2}\big) \\ &\geq (1-\tfrac{1}{N})\cdot \lambda_{\min}(WW')\cdot s_{\min}(H^{-1/2}AA'H^{-1/2}) \\ &= (1-\tfrac{1}{N})\cdot n\lambda_{\min}(\Sigma_W)\cdot\lambda_{\min}(\Sigma_A) \geq (1-\tfrac{1}{N})c_1^2 n, \end{aligned}$$

where the last inequality is from the assumption (6); in the first and second inequalities, we have used a result in linear algebra: for a matrix A and a positive definite matrix B, $s_{\min}(ABA') \geq \lambda_{\min}(B)\cdot s_{\min}(AA')$.[8] Similarly, we find that $\lambda_{\max}(G_0) \leq (1/c_1)n\lambda_{\max}(\Sigma_W)\lambda_{\max}(\Sigma_A) \leq Cn$. Then, the first claim holds.

Next, consider Lemma 3.3. Denote by Ξ_j' the j-th row of $\Xi = [\xi_1, \ldots, \xi_K]$. Recall that the matrix V is defined by $\Xi = M_0^{-1/2}AV$. As a result,

$$\Xi_j = [M_0(j,j)]^{-1/2}(Va_j),$$

where a_j' is the j-th row of A. First, by (49), we have $c_1 h_j \leq M_0(j,j) \leq h_j$. Second, by Lemma 6.1, $(VV')^{-1} = A'M_0^{-1}A$; so, $\|V\|^2 = \lambda_{\min}^{-1}(A'M_0^{-1}A) \leq \lambda_{\min}^{-1}(A'H^{-1}A) \leq c_1^{-1}$, where the last inequality is due to (6). Last, $\|a_j\| \leq \|a_j\|_1 = h_j$. Combing these results, we obtain:

$$\|\Xi_j\| \leq \frac{\|V\|\|a_j\|}{\sqrt{M_0(j,j)}} \leq \frac{(1/\sqrt{c_1})\cdot h_j}{\sqrt{c_1 h_j}} = \frac{\sqrt{h_j}}{c_1}.$$

Then, it follows from the Cauchy-Schwarz inequality that $\sum_{\ell=1}^K |\xi_\ell(j)| = \|\Xi_j\|_1 \leq \sqrt{K}\|\Xi_j\| \leq C\sqrt{h_j}$. □

6.7. *Proof of Lemmas 3.4-3.5.* Write $Z = [z_1, \ldots, z_n] = [Z_1, \ldots, Z_p]'$. By elementary properties of multinomial distributions, $\mathrm{Cov}(z_i) = N^{-1}\mathrm{diag}(d_i^0) - N^{-1}d_i^0(d_i^0)'$. As a result,

$$E[ZZ'] = \sum_{i=1}^n \mathrm{Cov}(z_i) = \frac{n}{N}M_0 - \frac{1}{N}D_0D_0'.$$

[8] Proof: For any vector v, we have $v'ABA'v \geq \lambda_{\min}(B)\cdot\|A'v\|^2 = \lambda_{\min}(B)\cdot(v'AA'v)$.

Then, we can write $G - G_0 = E_1 + E_2 + E_3 + E_4$, where

$$\begin{aligned}
E_1 &= \frac{n}{N} M^{-1/2}(M_0 - M)M^{-1/2},\\
E_2 &= M^{-1/2}(D_0 Z' + ZD_0')M^{-1/2},\\
E_3 &= M^{-1/2}(ZZ' - E[ZZ'])M^{-1/2},\\
E_4 &= (1 - \frac{1}{N})\big(M^{-1/2} D_0 D_0' M^{-1/2} - M_0^{-1/2} D_0 D_0' M_0^{-1/2}\big).
\end{aligned}$$

Consider E_1. By Lemma 6.3, with probability $1 - o(n^{-3})$, $|M(j,j) - M_0(j,j)| \le C(Nn)^{-1/2}\sqrt{h_j \log(n)}$ for all $j = 1, \ldots, p$. Moreover, by (49), $c_1 h_j \le M_0(j,j) \le h_j$. Since $h_j \ge h_{\min} \gg (Nn)^{-1}\log(n)$, the above suggests that $|M(j,j) - M_0(j,j)| \ll M_0(j,j)$; in particular, $M(j,j) \ge M_0(j,j)/2$. As a result, with probability $1 - o(n^{-3})$, for all $1 \le j \le p$,

$$\|e_j' E_1\| \le \frac{n}{N} \frac{|M(j,j) - M_0(j,j)|}{M_0(j,j)/2} \le \frac{C\sqrt{n\log(n)}}{N\sqrt{Nh_j}}. \tag{27}$$

Also, with probability $1 - o(n^{-3})$,

$$\|E_1\| \le \frac{n}{N} \max_{1\le j\le p}\Big\{\frac{|M(j,j) - M_0(j,j)|}{M_0(j,j)/2}\Big\} \le \frac{C\sqrt{n\log(n)}}{N\sqrt{Nh_{\min}}}. \tag{28}$$

Consider E_2. Denote by W_k' the k-th row of W, and recall that A_k is the k-th column of A, $1 \le k \le K$. Then, $D_0 = \sum_{k=1}^K A_k W_k'$. It follows that

$$E_2 = \sum_{k=1}^K \big[(M^{-1/2}A_k)(M^{-1/2}ZW_k)' + (M^{-1/2}ZW_k)(M^{-1/2}A_k)'\big].$$

As a result, with probability $1 - o(n^{-3})$,

$$\|E_2\| \le \sum_{k=1}^K 2\|M^{-1/2}A_k\| \cdot \|M^{-1/2}ZW_k\| \le C\sum_{k=1}^K \|H^{-1/2}A_k\| \cdot \|M_0^{-1/2}ZW_k\|,$$

where the last inequality is because $M_0(j,j) \ge c_1 h_j$ and $M(j,j) \ge M_0(j,j)/2$ with probability $1 - o(n^{-3})$. By Lemma 6.4, $\|M_0^{-1/2}ZW_k\| \le CN^{-1/2}\sqrt{np\log(n)}$. Moreover, $\sum_{k=1}^K \|H^{-1/2}A_k\|^2 = \sum_{k=1}^K \sum_{j=1}^p h_j^{-1} A_k^2(j) \le \sum_{k=1}^K \sum_{j=1}^p A_k(j) = K$. It then follows from the Cauchy-Schwarz inequality that $\sum_{k=1}^K \|H^{-1/2}A_k\| \le K$. As a result, with probability $1 - o(n^{-3})$,

$$\|E_2\| \le CN^{-1/2}\sqrt{np\log(n)}. \tag{29}$$

In addition, with probability $1-o(n^{-3})$,

$$\begin{aligned}\|e_j'E_2\| &\le \sum_{k=1}^K \frac{A_k(j)}{\sqrt{M(j,j)}}\|M^{-1/2}ZW_k\| + \sum_{k=1}^K \frac{|Z_j'W_k|}{\sqrt{M(j,j)}}\|M^{-1/2}A_k\| \\ &\le C\sqrt{h_j}\max_{1\le k\le K}\|M_0^{-1/2}ZW_k\| + \frac{C}{\sqrt{h_j}}\max_{1\le k\le K}|Z_j'W_k| \\ &\le CN^{-1/2}\sqrt{nph_j\log(n)} + CN^{-1/2}\sqrt{n\log(n)} \\ &\le C\sqrt{\frac{n\log(n)}{N}}\big(1+\sqrt{ph_j}\big), \qquad (30)\end{aligned}$$

where the second inequality is due to that $M(j,j)\ge M_0(j,j)/2\ge c_1h_j/2$, $\sum_{k=1}^K A_k(j)=h_j$ and $\sum_{k=1}^K\|M^{-1/2}A_k\|\le\sqrt{2/c_1}\sum_{k=1}^K\|H^{-1/2}A_k\|\le K\sqrt{2/c_1}$, and the third inequality follows from Lemma 6.4.

Consider E_3. We have seen that $\|M^{-1/2}M_0^{1/2}\|\le 2$ with probability $1-o(n^{-3})$. Combining it with Lemma 6.6 gives: with probability $1-o(n^{-3})$,

$$\|E_3\|\le 2\|M_0^{-1/2}(ZZ'-E[ZZ'])M_0^{-1/2}\|\le C\Big(\frac{1}{N}+\frac{p}{N^2h_{\min}}\Big)\sqrt{np}. \qquad (31)$$

Furthermore, by Lemma 6.5, with probability $1-o(n^{-3})$, for all $1\le j,\ell\le p$,

$$|E_3(j,\ell)| = \frac{|Z_j'Z_\ell - E[Z_j'Z_\ell]|}{\sqrt{M(j,j)M(\ell,\ell)}} \le \frac{C}{\sqrt{h_jh_\ell}}\cdot\frac{\sqrt{nh_jh_\ell\log(n)}}{N} \le \frac{C\sqrt{n\log(n)}}{N}.$$

It follows that with probability $1-o(n^{-3})$.

$$\|e_j'E_3\|\le CN^{-1}\sqrt{np\log(n)}. \qquad (32)$$

Consider E_4. Since $D_0=\sum_{k=1}^K A_kW_k'$,

$$\begin{aligned}E_4 &= (1-\frac{1}{N})\sum_{k,\ell=1}^K (W_k'W_\ell)\big(M^{-1/2}A_kA_\ell'M^{-1/2} - M_0^{-1/2}A_kA_\ell'M_0^{-1/2}\big) \\ &= (1-\frac{1}{N})\sum_{k,\ell=1}^K (W_k'W_\ell)\big[M^{-1/2}A_kA_\ell'(M^{-1/2}-M_0^{-1/2}) + (M^{-1/2}-M_0^{-1/2})A_kA_\ell'M_0^{-1/2}\big].\end{aligned}$$

In the proof of (29)-(30), we have seen that $\sum_{k=1}^K\|M^{-1/2}A_k\|\le 2\sum_{k=1}^K\|M_0^{-1/2}A_k\|\le C$. It follows that

$$\|E_4\|\le n\sum_{k,\ell=1}^K\big(\|M^{-1/2}A_k\|\|(M^{-1/2}-M_0^{-1/2})A_\ell\| + \|M_0^{-1/2}A_\ell\|\|(M^{-1/2}-M_0^{-1/2})A_k\|\big)$$

$$\leq CnK \cdot \max_{1\leq k\leq K} \|(M^{-1/2} - M_0^{-1/2})A_k\|.$$

By Lemma 6.3 and that $M(j,j) \geq M_0(j,j)/2 \geq c_1 h/2$, with probability $1-o(n^{-3})$, $|[M(j,j)]^{-1/2} - [M_0(j,j)]^{-1/2}| \leq h_j^{-1}(Nn)^{-1/2}\sqrt{\log(n)}$. So, with probability $1-o(n^{-3})$,

$$\|(M^{-1/2} - M_0^{-1/2})A_k\| \leq \frac{\sqrt{\log(n)}}{\sqrt{Nn}} \sqrt{\sum_{j=1}^{p} h_j^{-2} A_k^2(j)} \leq \frac{C\sqrt{p\log(n)}}{\sqrt{Nn}}.$$

Combining the above, with probability $1-o(n^{-3})$,

$$\|E_4\| \leq CN^{-1/2}\sqrt{np\log(n)}. \tag{33}$$

Moreover,

$$\begin{aligned}
\|e_j' E_4\| &\leq \frac{n}{\sqrt{M(j,j)}} \cdot \sum_{k,\ell=1}^{K} A_k(j)\|(M^{-1/2} - M_0^{-1/2})A_\ell\| \\
&\quad + n\Big|\frac{1}{\sqrt{M(j,j)}} - \frac{1}{\sqrt{M_0(j,j)}}\Big| \cdot \sum_{k,\ell=1}^{K} A_k(j)\|M_0^{-1/2}A_\ell\| \\
&\leq C\frac{n}{\sqrt{h_j}} \cdot h_j \cdot \frac{\sqrt{p\log(n)}}{\sqrt{Nn}} + Cn \cdot \frac{\sqrt{\log(n)}}{h_j\sqrt{Nn}} \cdot h_j \\
&\leq C\sqrt{\frac{n\log(n)}{N}}\big(1 + \sqrt{ph_j}\big). \qquad (34)
\end{aligned}$$

We now combine the results on E_1-E_4. By (27), (30), (32) and (34), with probability $1-o(n^{-3})$,

$$\begin{aligned}
\|e_j'(G - G_0)\| &\leq C\sqrt{\frac{n\log(n)}{N}}\Big(1 + \sqrt{ph_j} + \frac{1}{N\sqrt{h_j}} + \frac{\sqrt{p}}{\sqrt{N}}\Big) \\
&\leq C\sqrt{\frac{n\log(n)}{N}}\big(1 + \sqrt{ph_j}\big),
\end{aligned}$$

where the last inequality is due to that $h_j \geq h_{\min} \gg N^{-1}\log(n)$. This gives Lemma 3.4. By (28), (29), (31) and (33), with probability $1-o(n^{-3})$,

$$\begin{aligned}
\|G - G_0\| &\leq C\sqrt{\frac{np\log(n)}{N}}\Big(1 + \frac{1}{N\sqrt{ph_{\min}}} + \frac{1}{\sqrt{N}} + \frac{p}{h_{\min}N\sqrt{N}}\Big) \\
&\leq C\sqrt{\frac{np\log(n)}{N}}\Big(1 + \frac{p}{h_{\min}N^{3/2}}\Big),
\end{aligned}$$

where the last inequality is because $h_{\min} \gg N^{-1}\log(n)$ implies $\frac{1}{N\sqrt{ph_{\min}}} \ll \frac{1}{\sqrt{Np\log(n)}} = o(1)$. This gives Lemma 3.5. □

6.8. *Proof of Lemma 3.6.* Consider the first claim. Recall that $\hat{\Xi}_j'$ and Ξ_j' are the j-th row of $\hat{\Xi}$ and Ξ, respectively, $1 \le j \le p$. Write $\hat{\Xi}_j = (\hat{\xi}_1(j), \hat{b}_j')'$ and $\Xi_j = (\xi_1(j), b_j')'$. Then,

$$r_j = [\xi_1(j)]^{-1} b_j, \qquad \hat{r}_j = [\hat{\xi}_1(j)]^{-1}\hat{b}_j.$$

Let $\Omega = \mathrm{diag}(\omega_1, \ldots, \omega_K)$ be the minimizer in the definition of $\Delta_1(Z, D_0)$. Define $\Omega_1 = \mathrm{diag}(\omega_2, \ldots, \omega_K)$ and $\Omega^* = \omega_1^{-1}\Omega_1$. It is seen that

$$\begin{aligned}
\|\Omega^*\hat{r}_j - r_j\| &= \|\frac{1}{\hat{\xi}_1(j)}\Omega^*\hat{b}_j - \frac{1}{\xi_1(j)}b_j\| = \|\frac{1}{\omega_1\hat{\xi}_1(j)}\Omega_1\hat{b}_j - \frac{1}{\xi_1(j)}b_j\| \\
&= \Big\|\frac{1}{\omega_1\hat{\xi}_1(j)}(\Omega_1\hat{b}_j - b_j) - \frac{\omega_1\hat{\xi}_1(j) - \xi_1(j)}{\omega_1\hat{\xi}_1(j)}r_j\Big\| \\
&\le \frac{1}{|\omega_1\hat{\xi}_1(j)|}\big(\|\Omega_1\hat{b}_j - b_j\| + \|r_j\| \cdot |\omega_1\hat{\xi}_1(j) - \xi_1(j)|\big). \qquad (35)
\end{aligned}$$

By the definition of $\Delta_1(Z, D_0)$, $\|\Omega\hat{\Xi}_j - \Xi_j\| \le \sqrt{h_j}\Delta_1(Z, D_0)$. It follows that

$$\max\big\{|\omega_1\hat{\xi}_1(j) - \xi_1(j)|, \|\Omega_1\hat{b}_j - b_j\|\big\} \le \sqrt{h_j}\Delta_1(Z, D_0).$$

Additionally, by Lemma 6.2, $|\xi_1(j)| \ge C_5^{-1}\sqrt{h_j}$. It follows that $|\omega_1\hat{\xi}_1(j)| \ge |\xi_1(j)| - |\omega_1\hat{\xi}_1(j) - \xi_1(j)| \ge [C_5^{-1} - \Delta_1(Z, D_0)]\sqrt{h_j} \ge c\sqrt{h_j}$, for a constant $c > 0$. Last, by Lemma 6.2 again, $\|r_j\| \le C_6$. Plugging these results into (35), we find that

$$\|\Omega^*\hat{r}_j - r_j\| \le C\Delta_1(Z, D_0). \qquad (36)$$

Consider the second claim. By Lemma 6.2, for a constant $C > 1$, $C^{-1}\|\tilde{a}_i - \tilde{a}_j\| \le \|r_i - r_j\| \le C\|\tilde{a}_i - \tilde{a}_j\|$, for all $1 \le i, j \le p$. This allows us to translate the assumption (9) to conditions on $\{r_j\}_{j=1}^p$. First, since $\tilde{a}_j = e_k$ for j being an anchor word of topic k, the first part of (9) implies

$$\min_{j \in \mathcal{C}} \min_{1 \le k \le K} \|r_j - v_k^*\| \ge c, \qquad (37)$$

for a constant $c > 0$. Second, we study the k-means loss associated with the r_j's of non-anchor words. From the second part of (9), there exist $\eta_1^*, \ldots, \eta_{L_0}^*$ such that $\sum_{j \in \mathcal{C}}\{\min_{1 \le \ell \le L_0} \|\tilde{a}_j - \eta_\ell^*\|^2\} \le m_p/\log(n)$. Define a mapping $\mathcal{R}$ which maps a weight vector $\tilde{a} \in \mathbb{R}^K$ to a vector $r \in \mathbb{R}^{K-1}$ as follows: (Here $\circ$ denotes the entry-wise product and V_1 is the first column of V)

$$\tilde{a} \quad \mapsto \quad r \equiv \mathcal{R}\tilde{a} = [v_1^*, \ldots, v_K^*]\pi, \quad \text{where} \quad \pi = \frac{V_1 \circ \tilde{a}}{\|V_1 \circ \tilde{a}\|_1}.$$

From the proof of Lemma 6.2, we find that (i) $\mathcal{R}\tilde{a}_j = r_j$ for all $1 \le j \le p$, and (ii) for any two weight vectors $\tilde{a}$ and $\tilde{b}$, $C^{-1}\|\tilde{a}-\tilde{b}\| \le \|\mathcal{R}\tilde{a}-\mathcal{R}\tilde{b}\| \le C\|\tilde{a}-\tilde{b}\|$. Let $\theta^*_\ell = \mathcal{R}\eta^*_\ell$, $1 \le \ell \le L_0$. It is seen that

$$\begin{aligned}\sum_{j\in\mathcal{C}}\big\{\min_{1\le\ell\le L_0}\|r_j-\theta^*_\ell\|^2\big\} &= \sum_{j\in\mathcal{C}}\big\{\min_{1\le\ell\le L}\|\mathcal{R}\tilde{a}_j-\mathcal{R}\eta^*_\ell\|^2\big\}\\ &\le C\sum_{j\in\mathcal{C}}\big\{\min_{1\le\ell\le L}\|\tilde{a}_j-\eta^*_\ell\|^2\big\} \le C\frac{m_p}{\log(n)}.\end{aligned}$$

Write $\Theta = \{v^*_1,\ldots,v^*_K,\theta^*_1,\ldots,\theta^*_{L_0}\}$. Combining the above with the fact that $r_j = v^*_k$ for j being an anchor word of topic k, we have

$$\sum_{j=1}^{p}\big\{\min_{\theta\in\Theta}\|r_j-\theta\|^2\big\} \le Cm_p/\log(n). \tag{38}$$

We are ready to show the claim. Recall that the VH-1 step runs k-means on $\{\hat{r}_j\}_{j=1}^p$ and $\hat{\theta}^*_1,\ldots,\hat{\theta}^*_L$ are the resultant local centers. Introduce

$$\begin{aligned}\mathcal{B}^*_k &= \big\{\theta\in\mathbb{R}^{K-1}: \|\theta-(\Omega^*)^{-1}v^*_k\| \le C\Delta_1(Z,D_0)\big\}, \quad 1\le k\le K\\ \mathcal{C}^* &= \big\{\theta\in\mathbb{R}^{K-1}: \min_{j\in\mathcal{C}}\|\theta-(\Omega^*)^{-1}r_j\| \le C\Delta_1(Z,D_0)\big\}.\end{aligned}$$

We claim:

(i) The $(K+1)$ sets are disjoint from each other, and the distance between any two of them is at least $c/2$.
(ii) All local centers $\hat{\theta}^*_1,\ldots,\hat{\theta}^*_L$ are contained in the union of these sets.
(iii) Each of $\mathcal{B}^*_1,\ldots,\mathcal{B}^*_K$ contains at least one local center $\hat{\theta}^*_\ell$.

Once (i)-(iii) are true, in the VH-2 step, the algorithm has to select a single $\hat{\theta}^*_\ell$ from each of $\mathcal{B}^*_1,\ldots,\mathcal{B}^*_K$ (no matter whether we use the combinatorial search or the greedy search). The claim then follows.

It remains to prove (i)-(iii). Here, (i) follows from (37) and that Ω^* is an orthogonal matrix. To see (ii), we note that by (36) all $\hat{r}_j$'s are contained in the union of $\mathcal{B}^*_1,\ldots,\mathcal{B}^*_K$ and $\mathcal{C}$, so the k-means algorithm will also place all local centers in the union of these sets.

Below, we show (iii). Using the universal inequality $(a+b)^2 \le 2a^2+2b^2$, we have

$$\sum_{j=1}^{p}\big\{\min_{1\le\ell\le L}\|r_j-\Omega^*\hat{\theta}^*_\ell\|^2\big\} \le 2\sum_{j=1}^{p}\big\{\min_{1\le\ell\le L}\|\Omega^*\hat{r}_j-\Omega^*\hat{\theta}^*_\ell\|^2\big\} + 2\sum_{j=1}^{p}\|\Omega^*\hat{r}_j-r_j\|^2$$

$$
\begin{aligned}
&\le 2\sum_{j=1}^{p}\big\{\min_{1\le\ell\le L}\|\hat{r}_j-\hat{\theta}^*_\ell\|^2\big\}+Cp\Delta_1^2(Z,D_0)\\
&\le 2\sum_{j=1}^{p}\big\{\min_{\theta\in\Theta}\|\hat{r}_j-(\Omega^*)^{-1}\theta\|^2\big\}+Cp\Delta_1^2(Z,D_0)\\
&= 2\sum_{j=1}^{p}\big\{\min_{\theta\in\Theta}\|\Omega^*\hat{r}_j-\theta\|^2\big\}+Cp\Delta_1^2(Z,D_0)\\
&\le 4\sum_{j=1}^{p}\big\{\min_{\theta\in\Theta}\|r_j-\theta\|^2\big\}+4\sum_{j=1}^{p}\|\Omega^*\hat{r}_j-r_j\|^2+Cp\Delta_1^2(Z,D_0)\\
&\le Cm_p/\log(n)+Cp\Delta_1^2(Z,D_0).
\end{aligned}
$$

Here, the third line is because $\{(\Omega^*)^{-1}\theta:\theta\in\Theta\}$ serves as a feasible solution for the k-means problem, the fifth line uses the inequality $(a+b)^2\le 2a^2+2b^2$ again, and the second and last lines are from (36) and (38). Since $m_p/p\ge\Delta_1^2(Z,D_0)\log(n)$ by the assumptions, we have

$$
\sum_{j=1}^{p}\big\{\min_{1\le\ell\le L}\|r_j-\Omega^*\hat{\theta}^*_\ell\|^2\big\}\le Cm_p/\log(n). \tag{39}
$$

Now, suppose (iii) does not hold, i.e., there exists a $\mathcal{B}^*_k$ that contains no local centers. By (i)-(ii), all local centers are at least a distance $c/2$ from $(\Omega^*)^{-1}v^*_k$. As a result, for any anchor word j of topic k, all local centers are at least a distance $c/2$ from $(\Omega^*)^{-1}r_j$. It follows that $\min_{1\le\ell\le L}\|r_j-\Omega^*\hat{\theta}^*_\ell\|=\min_{1\le\ell\le L}\|(\Omega^*)^{-1}r_j-\hat{\theta}^*_\ell\|\ge c/2$. Since each topic has at least m_p anchor words,

$$
\sum_{j=1}^{p}\big\{\min_{1\le\ell\le L}\|r_j-\Omega^*\hat{\theta}^*_\ell\|^2\big\}\ge(c/2)^2\cdot m_p.
$$

This yields a contradiction to (39). So (iii) must hold. □

6.9. *Proof of Lemma 3.7.* For notation simplicity, in the proof below, we omit the permutation $\kappa(\cdot)$ in the definition of Err_{VH}. From the definitions of $\Delta_1(Z,D_0)$, $\Delta_2(Z,D_0)$ and Err_{VH}, there exist $\Omega\in\mathcal{O}_K$ and $\Omega^*\in\mathcal{O}_{K-1}$ such that for all $1\le j\le p, 1\le k\le K$,

$$
\begin{cases}
\|\Omega\hat{\Xi}_j-\Xi_j\|\le\Delta_1(Z,D_0)\cdot\sqrt{h_j},\\
\|M(j,j)-M_0(j,j)\|\le\Delta_2(Z,D_0)\cdot h_j,\\
\|\Omega^*\hat{v}^*_k-v^*_k\|\le Err_{VH}.
\end{cases}
$$

Moreover, from the proof of Lemma 3.6, Ω^* is uniquely determined by Ω: letting $\Omega = \mathrm{diag}(\omega_1, \ldots, \omega_K)$, we have $\Omega^* = \omega_1^{-1}\mathrm{diag}(\omega_2, \ldots, \omega_K)$.

First, we consider the step of weight matrix estimation. For each $1 \le j \le p$, $\hat{\pi}_j$ is obtained by truncating and renormalizing $\hat{\pi}_j^*$, where $\hat{\pi}_j^*$ solves the linear equation

$$\begin{pmatrix} 1 & \ldots & 1 \\ \hat{v}_1^* & \ldots & \hat{v}_K^* \end{pmatrix} \hat{\pi}_j^* = \begin{pmatrix} 1 \\ \hat{r}_j \end{pmatrix} \iff \begin{pmatrix} 1 & \ldots & 1 \\ \Omega^*\hat{v}_1^* & \ldots & \Omega^*\hat{v}_K^* \end{pmatrix} \hat{\pi}_j^* = \begin{pmatrix} 1 \\ \Omega^*\hat{r}_j \end{pmatrix}.$$

It follows that

$$\hat{\pi}_j^* = \hat{Q}^{-1} \begin{pmatrix} 1 \\ \Omega^*\hat{r}_j \end{pmatrix}, \quad \text{where } \hat{Q} = \begin{pmatrix} 1 & \ldots & 1 \\ \Omega^*\hat{v}_1^* & \ldots & \Omega^*\hat{v}_K^* \end{pmatrix}.$$

Moreover, by Lemma 1.1, π_j is a PMF which satisfies that $\sum_{k=1}^K \pi_j(k) v_k^* = r_j$. Similarly, we have

$$\pi_j = Q^{-1} \begin{pmatrix} 1 \\ r_j \end{pmatrix}, \quad \text{where } Q = \begin{pmatrix} 1 & \ldots & 1 \\ v_1^* & \ldots & v_K^* \end{pmatrix}.$$

Consequently,

$$\|\hat{\pi}_j^* - \pi_j\| \le \|\hat{Q}^{-1}\| \|\Omega^*\hat{r}_j - r_j\| + \|\hat{Q}^{-1} - Q^{-1}\| \|r_j\|. \tag{40}$$

Noting that $Q' = [\mathrm{diag}(V_1)]^{-1}V$, we have $\|Q^{-1}\|^2 = \|(Q'Q)^{-1}\|^2 \le (\max_k |V_1(k)|)^2 \cdot \|(VV')^{-1}\|$. By Lemma 6.1, $(VV')^{-1} = A'M_0^{-1}A$; moreover, by (49), $\|A'M_0^{-1}A\| \le c_1^{-1}\|A'H^{-1}A\|$; recalling that a_j' is the j-th row of A, we find that $\|A'H^{-1}A\| \le \|A'H^{-1}A\|_1 = \max_k \sum_{\ell=1}^K \sum_{j=1}^p \|a_j\|_1^{-1} a_j(k) a_j(\ell) \le \max_k \sum_{\ell=1}^K \sum_{j=1}^p a_j(\ell) = K$. Furthermore, by Lemma 6.1 again, $C^{-1} \le |V_1(k)| \le C$ for all $1 \le k \le K$. Combining the above gives that

$$\|Q^{-1}\| \le C.$$

Additionally, it is easy to see that $\|\hat{Q} - Q\| \le \|\hat{Q} - Q\|_1 \le \sqrt{K} \max_k \|\Omega^*\hat{v}_k^* - v_k^*\|$; as a result, $\|\hat{Q}^{-1} - Q^{-1}\| \le \|\hat{Q}^{-1}\| \|Q^{-1}\| \|\hat{Q} - Q\| \le C \max_k \|\Omega^*\hat{v}_k^* - v_k^*\|$. Moreover, by Lemma 6.2, $\|r_j\| \le C$. Combining the above, we find that

$$\begin{aligned} \|\hat{\pi}_j^* - \pi_j\| &\le C\big(\|\Omega^*\hat{r}_j - r_j\| + \max_{1\le k\le K} \|\Omega^*\hat{v}_k^* - v_k^*\|\big) \\ &\le C(\|\Omega^*\hat{r}_j - r_j\| + Err_{VH}). \end{aligned} \tag{41}$$

By definition, $\hat{\pi}_j = \tilde{\pi}_j^*/\|\tilde{\pi}_j^*\|_1$, where $\tilde{\pi}_j^*(k) = \max\{\hat{\pi}_j^*(k), 0\}$. It is seen that

$$\|\hat{\pi}_j - \pi_j\|_1 \le \|\hat{\pi}_j - \tilde{\pi}_j^*\|_1 + \|\tilde{\pi}_j^* - \pi_j\|_1$$

$$= \|(1 - \|\tilde{\pi}_j^*\|_1)\hat{\pi}_j\|_1 + \|\tilde{\pi}_j^* - \pi_j\|_1$$
$$= |1 - \|\tilde{\pi}_j^*\|_1| + \|\tilde{\pi}_j^* - \pi_j\|_1.$$

Using the triangle inequality, we have $|1-\|\tilde{\pi}_j^*\|_1| \le \|\tilde{\pi}_j^*-\pi_j\|_1$. Furthermore, since all entries of π_j are nonnegative, $\|\tilde{\pi}_j^* - \pi_j\|_1 \le \|\hat{\pi}_j^* - \pi_j\|_1 \le \sqrt{K}\|\hat{\pi}_j^* - \pi_j\|$. As a result,

$$\|\hat{\pi}_j - \pi_j\|_1 \le 2\sqrt{K}\|\hat{\pi}_j^* - \pi_j\| \le C(\|\Omega^*\hat{r}_j - r_j\| + Err_{VH}). \tag{42}$$

Next, consider the step of topic matrix estimation. Given $\hat{\Pi} = [\hat{\pi}_1, \ldots, \hat{\pi}_p]'$, $\hat{A}$ is obtained by re-normalizing each column of $\hat{A}^* = M^{1/2} \cdot \mathrm{diag}(\hat{\xi}_1) \cdot \hat{\Pi}$ to have a unit ℓ_1 norm. Since we always flip the sign of $\hat{\xi}_1$ such that the entries of $\hat{A}^*$ are non-negative, we write $\hat{A}^* = M^{1/2} \cdot \mathrm{diag}(\omega_1\hat{\xi}_1) \cdot \hat{\Pi}$, where $\omega_1 \in \{\pm 1\}$ is the first diagonal of Ω. Introduce $A^* = M_0^{1/2} \cdot \mathrm{diag}(\xi_1) \cdot \Pi$. By Lemma 1.1,

$$A^* = A \cdot \mathrm{diag}(V_1). \tag{43}$$

Denote by $(\hat{a}_j^*)'$ and $(a_j^*)'$ the j-th row of $\hat{A}^*$ and A^*, respectively, $1 \le j \le p$. Then,

$$\begin{aligned}\|\hat{a}_j^* - a_j^*\|_1 &= \|\sqrt{M(j,j)}\omega_1\hat{\xi}_1(j)\hat{\pi}_j - \sqrt{M_0(j,j)}\xi_1(j)\pi_j\|_1 \\ &\le \sqrt{M(j,j)} \cdot |\omega_1\hat{\xi}_1(j)| \cdot \|\hat{\pi}_j - \pi_j\|_1 + \sqrt{M(j,j)}\|\pi_j\|_1 \cdot |\omega_1\hat{\xi}_1(j) - \xi_1(j)| \\ &\quad + |\xi_1(j)|\|\pi_j\|_1 \cdot |\sqrt{M(j,j)} - \sqrt{M_0(j,j)}|.\end{aligned}$$

First, $|\omega_1\hat{\xi}_1(j) - \xi_1(j)| \le \|\Omega\hat{\Xi}_j - \Xi_j\| \le \sqrt{h_j}\Delta_1(Z, D_0)$. Second, by Lemma 6.2, $|\xi_1(j)| \le C\sqrt{h_j}$; furthermore, $|\omega_1\hat{\xi}_1(j)| \le 2|\xi_1(j)| \le C\sqrt{h_j}$. Third, by (49) and the definition of $\Delta_2(Z, D_0)$, $|\sqrt{M(j,j)} - \sqrt{M_0(j,j)}| \le C\sqrt{h_j} \cdot \Delta_2(Z, D_0)$ and $M(j,j) \le 2M_0(j,j) \le Ch_j$. As a result,

$$\begin{aligned}\|\hat{a}_j^* - a_j^*\|_1 &\le C\sqrt{h_j} \cdot |\omega_1\hat{\xi}_1(j)| \cdot \|\hat{\pi}_j - \pi_j\|_1 + Ch_j \cdot \Delta(Z, D_0) \\ &\le C\sqrt{h_j} \cdot |\omega_1\hat{\xi}_1(j)| \cdot (\|\Omega^*\hat{r}_j - r_j\| + Err_{VH}) + Ch_j \cdot \Delta(Z, D_0) \\ &\le C\sqrt{h_j} \cdot |\omega_1\hat{\xi}_1(j)| \cdot \|\Omega^*\hat{r}_j - r_j\| + Ch_j \cdot [\Delta(Z, D_0) + Err_{VH}], \qquad (44)\end{aligned}$$

where we have used (42) in the second inequality. Below, we bound $\|\Omega^*\hat{r}_j - r_j\|$. Write $\hat{\Xi}_j = (\hat{\xi}_1(j), \hat{b}_j')'$ and $\Xi_j = (\xi_1(j), b_j')'$. By definition, $\hat{r}_j = [\hat{\xi}_1(j)]^{-1}\hat{b}_j$ and $r_j = [\xi_1(j)]^{-1}b_j$. Direct calculations give

$$\omega_1\hat{\xi}_1(j) \cdot (\Omega^*\hat{r}_j - r_j) = (\Omega^*\hat{b}_j - b_j) - [\omega_1\hat{\xi}_1(j) - \xi_1(j)]r_j.$$

Note that $\max\{|\omega_1\hat{\xi}_1(j)-\xi_1(j)|, \|\Omega^*\hat{b}_j-b_j\|\} \le \|\Omega\hat{\Xi}_j-\Xi_j\| \le \Delta_1(Z,D_0)\sqrt{h_j}$. In addition, by Lemma 6.2, $\|r_j\| \le C$. It follows that

$$|\omega_1\xi_1(j)| \cdot \|\Omega^*\hat{r}_j - r_j\| \le C\sqrt{h_j} \cdot \Delta_1(Z,D_0).$$

Plugging it into (44) gives

$$\|\hat{a}_j^* - a_j^*\|_1 \le Ch_j \cdot [\Delta(Z,D_0) + Err_{VH}]. \tag{45}$$

Denote by $\hat{A}_k$ and $\hat{A}_k^*$ the k-th column of $\hat{A}$ and $\hat{A}^* = [\hat{a}_1^*, \ldots, \hat{a}_p^*]'$, respectively, $1 \le k \le K$. Then, $\hat{A}_k = \|\hat{A}_k^*\|_1^{-1}\hat{A}_k^*$. Combining it with (43) gives

$$\hat{a}_j(k) = \|\hat{A}_k^*\|_1^{-1} \cdot \hat{a}_j^*(k), \qquad a_j(k) = [V_1(k)]^{-1} \cdot a_j^*(k).$$

So, $|\hat{a}_j(k) - a_j(k)| \le \frac{1}{\|\hat{A}_k^*\|_1}|\hat{a}_j^*(k) - a_j^*(k)| + \frac{|\|\hat{A}_k^*\|_1 - V_1(k)|}{\|\hat{A}_k^*\|_1}|a_j(k)|$. Taking the sum over $k = 1, \ldots, K$ on both sides, we find that

$$\begin{aligned}\|\hat{a}_j - a_j\|_1 &\le \frac{1}{\|\hat{A}_k^*\|_1}\|\hat{a}_j^* - a_j^*\|_1 + \frac{|\|\hat{A}_k^*\|_1 - V_1(k)|}{\|\hat{A}_k^*\|_1}\|a_j\|_1 \\ &\le \frac{Ch_j}{\|\hat{A}_k^*\|_1}\Big[\Delta(Z,D_0) + Err_{VH} + |\|\hat{A}_k^*\|_1 - V_1(k)|\Big], \end{aligned} \tag{46}$$

where the last inequality is from (45) and that $\|a_j\|_1 = h_j$. Below, we bound $|\|\hat{A}_k^*\|_1 - V_1(k)|$. From (43), $\|A_k^*\|_1 = \|A_k\|_1 \cdot V_1(k) = V_1(k)$. Then, $|\|\hat{A}_k^*\|_1 - V_1(k)| = |\|\hat{A}_k^*\|_1 - \|A_k^*\|_1| \le \|\hat{A}_k^* - A_k^*\|_1 \le \sum_{j=1}^p |\hat{a}_j^*(k) - a_j^*(k)| \le \sum_{j=1}^p \|\hat{a}_j^* - a_j^*\|_1$. We then apply (45) again and use the fact that $\sum_{j=1}^p h_j = K$. It yields

$$|\|\hat{A}_k^*\|_1 - V_1(k)| \le C[\Delta(Z,D_0) + Err_{VH}].$$

Moreover, by Lemma 6.1, $V_1(k) \ge C^{-1}$; it follows that $\|\hat{A}_k^*\|_1 \ge V_1(k)/2 \ge C^{-1}$. Combining these results with (46) gives

$$\|\hat{a}_j - a_j\|_1 \le Ch_j \cdot [\Delta(Z,D_0) + Err_{VH}]. \tag{47}$$

The claim then follows (note that $h_j = \|a_j\|_1$). □

7. Supplementary proofs.

7.1. *Proof of Lemma 6.1.* Consider the first claim. Note that $M_0^{-1/2}D_0$ has a full column rank K. Let

$$M_0^{-1/2}D_0 = \Xi\Lambda B'$$

be the Singular Value Decomposition of $M_0^{-1/2}D_0$, where $\Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_K)$ contains the singular values and $B \in \mathbb{R}^{n,K}$ contains the right singular vectors; note that $\Xi'\Xi = B'B = I_K$. It is seen that

$$\Xi = (\Xi\Lambda B')B\Lambda^{-1} = M_0^{-1/2}D_0B\Lambda^{-1} = M_0^{-1/2}A(WB\Lambda^{-1}).$$

By letting $V = WB\Lambda^{-1}$, we have $\Xi = AV$; i.e., such a V exists. Furthermore, for any V such that $\Xi = M_0^{-1/2}AV$, we have $\Xi'M_0^{-1/2}AV = \Xi'\Xi = I_K$. This implies that V is the inverse of $(\Xi'M_0^{-1/2}A)$, so V is unique and non-singular. Last, we plug $\Xi = M_0^{-1/2}AV$ into $\Xi'\Xi = I_K$; it yields $I_K = V'A'M_0^{-1}AV$. Multiplying both sides of this equation by V from the left and by V' from the right, we obtain:

$$VV' = (VV')A'M_0^{-1}A(VV').$$

This proves that $VV' = (A'M_0^{-1}A)^{-1}$.

Consider the second claim. We first show that

$$|V_1(k)| \le C, \qquad \text{for } 1 \le k \le K. \tag{48}$$

We aim to use the fact that $VV' = (A'M_0^{-1}A)^{-1}$, so the key is to study the diagonal matrix M_0. Note that $M_0(j,j) = \frac{1}{n}\sum_{i=1}^n[\sum_{k=1}^K A_k(j)w_i(k)] = \sum_{k=1}^K A_k(j)[\frac{1}{n}\sum_{i=1}^n w_i(k)]$. Since $w_i(k) \le 1$, we have $M_0(j,j) \le \sum_{k=1}^K A_k(j) = h_j$. At the same time, $\frac{1}{n}\sum_{i=1}^n w_i(k) \ge \frac{1}{n}\sum_{i=1}^n w_i^2(k) = \Sigma_W(k,k)$, and it follows from the assumption (6) that $\Sigma_W(k,k) \ge c_1$; consequently, $M_0(j,j) \ge c_1\sum_{k=1}^K A_k(j) = c_1h_j$. In summary,

$$c_1h_j \le M_0(j,j) \le h_j, \qquad \text{for } 1 \le j \le p. \tag{49}$$

Recall the matrix $H = \mathrm{diag}(h_1, \ldots, h_p)$. By (49), $A'(M_0^{-1} - H^{-1})A$ is positive semi-definite, which implies $\lambda_{\min}(A'M_0^{-1}A) \ge \lambda_{\min}(A'H^{-1}A)$; similarly, $\lambda_{\max}(A'M_0^{-1}A) \le c_1^{-1}\lambda_{\max}(A'H^{-1}A)$. Note that $A'H^{-1}A = \Sigma_A$. By the assumption (6), $\lambda_{\min}(\Sigma_A) \ge c_1$; also, using the fact that the column sums of A are equal to 1, we have $\lambda_{\max}(\Sigma_A) \le \|\Sigma_A\|_1 = 1$. Combining the above gives

$$c_1 \le \lambda_{\min}(A'M_0^{-1}A) \le \lambda_{\max}(A'M_0^{-1}A) \le c_1^{-1}. \tag{50}$$

In the first claim, we have seen that $VV' = (A'M_0^{-1}A)^{-1}$. So, (50) yields:

$$c_1 \le \lambda_{\min}(VV') \le \lambda_{\max}(VV') \le c_1^{-1}. \tag{51}$$

Observing that $\sum_{\ell=1}^K V_\ell^2(k)$ is the k-th diagonal of VV', we obtain (48).

Next, we show that for a constant $c > 0$, up to a multiple of ± 1 on V_1,

$$V_1(k) \ge c, \qquad \text{for } 1 \le k \le K. \tag{52}$$

Let $\eta_1 = \text{sign}(V_1(1)) \cdot \|V_1\|^{-1} V_1$. Since $\|V_1\|^2$ is the first diagonal of $V'V$, we have $\|V_1\|^2 \ge \lambda_{\min}(V'V) = \lambda_{\min}(VV') \ge c_1$, where the last inequality is due to (51). Therefore, to show (52), it suffices to show that

$$\liminf_{n\to\infty} \min_{1\le k\le K} \{\eta_1(k)\} \ge c. \tag{53}$$

Let $\lambda_1, \ldots, \lambda_K$ be the singular values of $M_0^{-1/2} D_0$. Then, $M_0^{-1/2} D_0 D_0' M_0^{-1/2} \xi_k = \lambda_k^2 \xi_k$, where $D_0 = AW$ and $\xi_k = M_0^{-1/2} A V_k$. Combining these facts gives $(M_0^{-1/2} AWW'A' M_0^{-1/2})(M_0^{-1/2} A V_k) = \lambda_k^2 (M_0^{-1/2} A V_k)$. Multiplying both sides by $(A'M_0^{-1}A)^{-1} A' M_0^{-1/2}$ from the left, we have

$$(WW'A'M_0^{-1}A) V_k = \lambda_k^2 V_k.$$

This means V_k is an eigenvector of the matrix $n\Sigma_W (A'M_0^{-1}A)$ associated with the eigenvalue λ_k^2. In particular,

$$\eta_1 \text{ is the unit-norm leading eigenvector of } \Theta = \Sigma_W (A'M_0^{-1}A). \tag{54}$$

Write $\eta_1 = \eta_1^{(n)}$ to indicate its dependence on n; similar for other quantities. Suppose (53) is not true. Then, there exists k and a subsequence $\{n_m\}_{m=1}^\infty$ such that $\lim_{m\to\infty} \eta_1^{(n_m)}(k) = 0$. Furthermore, the spectral norm of Σ_W is bounded (because each column of W is a weight vector), and the spectral norm of $A'M_0^{-1}A$ is also bounded (by (50)). Therefore, there exists a subsequence of $\{n_m\}_{m=1}^\infty$ such that Θ tends to a fixed matrix Θ_0; without loss of generality, we assume this subsequence is $\{n_m\}_{m=1}^\infty$ itself. The above implies

$$\lim_{m\to\infty} \eta_1^{(n_m)}(k) = 0, \qquad \lim_{m\to\infty} \Theta^{(n_m)} = \Theta_0.$$

From the assumption (7), the eigengap of Θ is bounded below by a constant $c_2 > 0$. Using the sine-theta theorem (Davis and Kahan, 1970), we find that when $\Theta^{(n_m)} \to \Theta_0$, up to a multiple of ± 1 on $\eta_1^{(n_m)}$,

$$\eta_1^{(n_m)} \to q_0, \qquad q_0 \text{ is the unit-norm leading eigenvector of } \Theta_0.$$

Combining the above gives

$$q_0(k) = 0. \tag{55}$$

We then study the matrix Θ_0. Write $\Theta = \Theta_1 + \Theta_2$, where $\Theta_1 = \Sigma_W(A'H^{-1}A)$ and $\Theta_2 = \Sigma_W A'(M_0^{-1} - H^{-1})A$. By (49), all entries of Θ_2 are non-negative. Moreover, the assumption (6) yields that all entries of $A'H^{-1}A$ are lower bounded by a constant $c_1 > 0$; as a result, all entries of Θ_1 are lower bounded by a positive constant. Combining the above, all entries of Θ are lower bounded by a positive constant, which implies:

$$\Theta_0 \text{ is a strictly positive matrix.} \tag{56}$$

By Perron's theorem (Horn and Johnson, 1985), the leading unit-norm eigenvector (up to ± 1) of a positive matrix has all positive entries. So (55) and (56) are contradicting with each other. This proves (53); then, (52) follows.

Consider the last three claims. The key is to study the matrix

$$Q \equiv \begin{pmatrix} 1 & \dots & 1 \\ v_1^* & \dots & v_K^* \end{pmatrix}.$$

From how $v_1^*, \dots, v_K^*$ are define, $Q' = [\mathrm{diag}(V_1)]^{-1} \cdot V$. So

$$Q'Q = [\mathrm{diag}(V_1)]^{-1} V V' [\mathrm{diag}(V_1)]^{-1}.$$

In the second claim, we have seen that the entries of V_1 are either all positive or all negative; also, $C^{-1} \le |V_1(k)| \le C$ for all $1 \le k \le K$. Combining this with (51) gives

$$C^{-1} \le \lambda_{\min}(Q'Q) \le \lambda_{\max}(Q'Q) \le C. \tag{57}$$

We first study $\|v_k^*\|$ and $\|v_k^* - v_\ell^*\|$. Note that

$$\begin{pmatrix} 1 \\ v_k^* \end{pmatrix} = Qe_k, \qquad e_k\text{: the } k\text{-th standard basis of } \mathbb{R}^K.$$

Therefore, $\|v_k^*\| \le \|Q\| \le C$, $\|v_k^* - v_\ell^*\| \le \|Q\| \cdot \|e_k - e_\ell\| \le \sqrt{2}\|Q\| \le C$, and $\|v_k^* - v_\ell^*\|^2 \ge \|e_k - e_\ell\|^2 \cdot \lambda_{\min}(Q'Q) \ge C^{-1}$.

We then study the simplex $\mathcal{S}_K^*$. By (57), Q is non-singular. Hence, there cannot be a non-zero vector b such $Qb = 0$; note that $Qb = 0$ is equivalent to that $\sum_{k=1}^K b(k) = 0$ and $\sum_{k=1}^K b(k)v_k^* = 0$. This means the vectors $v_1^*, \dots, v_K^*$ are affinely independent; so $\mathcal{S}_K^*$ is a non-degenerate simplex. Regarding the volume of $\mathcal{S}_K^*$, since $\|v_k^*\|$ are all bounded by a constant, the volume is also

upper bounded. To give a lower bound of the volume, we introduce $\bar{v}^* = K^{-1}\sum_{k=1}^K v_k^*$ and investigate the distance between $\bar{v}^*$ and the boundary of $\mathcal{S}_K^*$. For any point u in $\mathcal{S}_K^*$, it is a convex combination of the vertices, i.e., for a weight vector b, $u = \sum_{k=1}^K b(k)v_k^*$. It is seen that

$$\begin{pmatrix} 0 \\ u - \bar{v}^* \end{pmatrix} = Q(b - \bar{e}), \qquad \bar{e} = (1/K, 1/K, \ldots, 1/K)'.$$

If u is on the boundary of $\mathcal{S}_K^*$, then b has at least one zero entry, so $\|b-\bar{e}\| \geq 1/K$. Combining this with (57) gives $\|u-\bar{v}^*\| \geq [\lambda_{\min}(Q'Q)]^{1/2} \cdot \|b-\bar{e}\| \geq c$, for a constant $c > 0$. This means the distance from an inner point $\bar{v}^*$ to the boundary of the simplex is lower bounded by a constant. Then, the volume of the simplex is bounded below from zero. □

7.2. *Proof of Lemma 6.2.* Consider the first claim. From $\Xi = M_0^{-1/2}AV$, we have $\xi_1(j) = [M_0(j,j)]^{-1/2}a_j'V_1$ for $1 \leq j \leq p$. Note that a_j is a non-negative vector with $\|a_j\|_1 \neq 0$ and that all entries of V_1 are either all positive or all negative; so the entries of $a_j'V_1$ all have the same sign. Consequently, the entries of ξ_1 also have the same sign; this means we can choose the sign of ξ_1 so that all the entries are positive.

Assuming all entries of ξ_1 and V_1 are positive, we now give lower/upper bound of $\xi_1(j)$, for $1 \leq j \leq p$. Since $\xi_1(j) = [M_0(j,j)]^{-1/2}a_j'V_1$,

$$\xi_1(j) \geq [M_0(j,j)]^{-1/2}\|a_j\|_1 \min_{1\leq k\leq K} V_1(k).$$

By definition, $\|a_j\|_1 = h_j$. By (49), $M_0(j,j) \leq h_j$. By Lemma 6.1, $V_1(k) \geq C^{-1}$ for all $1 \leq k \leq K$. Combining the above gives

$$\xi_1(j) \geq C^{-1}\sqrt{h_j}.$$

Similarly, we can prove that $\xi_1(j) \leq C\sqrt{h_j}$.

Consider the second claim. Since each r_j is in the simplex $\mathcal{S}_K^*$, it follows that $\|r_j\| \leq \max_{1\leq k\leq K}\|v_k^*\|$; by Lemma 6.1, $\max_{1\leq k\leq K}\|v_k^*\| \leq C$. The claim then follows.

Consider the third claim. By Lemma 1.1, each r_j is a convex combination of $v_1^*, \ldots, v_K^*$, where the weight vector π_j is the j-th row of $\Pi = [\mathrm{diag}(\xi_1)]^{-1} \cdot M_0^{-1/2}A \cdot \mathrm{diag}(V_1)$. So

$$\begin{pmatrix} 0 \\ r_i - r_j \end{pmatrix} = Q(\pi_i - \pi_j), \qquad \text{where } Q = \begin{pmatrix} 1 & \ldots & 1 \\ v_1^* & \ldots & v_K^* \end{pmatrix}.$$

In (57), we have seen that $C^{-1} \leq \lambda_{\min}(Q'Q) \leq \lambda_{\max}(Q'Q) \leq C$. So,

$$C^{-1}\|\pi_i - \pi_j\| \leq \|r_i - r_j\| \leq C\|\pi_i - \pi_j\|.$$

To show the claim, it suffices to prove that

$$C^{-1}\|\tilde{a}_i - \tilde{a}_j\| \le \|\pi_i - \pi_j\| \le C\|\tilde{a}_i - \tilde{a}_j\|. \tag{58}$$

We now show (58). We assume the sign of ξ_1 is chosen such that all entries of ξ_1 and V_1 are positive. Since $\Pi = [\mathrm{diag}(\xi_1)]^{-1} \cdot M_0^{-1/2} A \cdot \mathrm{diag}(V_1)$,

$$\begin{aligned} \pi_j &= [\xi_1(j)]^{-1}[M_0(j,j)]^{-1/2} \cdot \mathrm{diag}(V_1) a_j \\ &= [\xi_1(j)]^{-1}[M_0(j,j)]^{-1/2} h_j \cdot \mathrm{diag}(V_1)\tilde{a}_j \\ &\propto (V_1 \circ \tilde{a}_j), \end{aligned} \tag{59}$$

where $\circ$ denotes the entry-wise product of two vectors. Noting that both π_j and $\tilde{a}_j$ are weight vectors, we have $\pi_j = (V_1 \circ \tilde{a}_j)/\|V_1 \circ \tilde{a}_j\|_1$. Therefore,

$$\pi_i - \pi_j = \frac{(V_1 \circ \tilde{a}_i)}{\|V_1 \circ \tilde{a}_i\|_1} - \frac{(V_1 \circ \tilde{a}_j)}{\|V_1 \circ \tilde{a}_j\|_1} = \frac{V_1 \circ (\tilde{a}_i - \tilde{a}_j)}{\|V_1 \circ \tilde{a}_i\|_1} + \frac{\|V_1 \circ \tilde{a}_j\|_1 - \|V_1 \circ \tilde{a}_i\|_1}{\|V_1 \circ \tilde{a}_i\|_1}\pi_j.$$

By the triangle inequality, $|\|V_1 \circ \tilde{a}_j\|_1 - \|V_1 \circ \tilde{a}_i\|_1| \le \|(V_1 \circ \tilde{a}_j) - (V_1 \circ \tilde{a}_i)\|_1 = \|V_1 \circ (\tilde{a}_i - \tilde{a}_j)\|_1$. Moreover, $\|\pi_j\|_1 = 1$. It follows that

$$\|\pi_i - \pi_j\|_1 \le 2\frac{\|V_1 \circ (\tilde{a}_i - \tilde{a}_j)\|_1}{\|V_1 \circ \tilde{a}_i\|_1}.$$

By Lemma 6.1, $C^{-1} \le V_1(k) \le C$ for all k. So $\|V_1 \circ (\tilde{a}_i - \tilde{a}_j)\|_1 \le C\|\tilde{a}_i - \tilde{a}_j\|_1$, and $\|V_1 \circ \tilde{a}_i\|_1 \ge C^{-1}$. It follows that

$$\|\pi_i - \pi_j\|_1 \le C\|\tilde{a}_i - \tilde{a}_j\|_1.$$

Using the Cauchy-Schwarz inequality, $\|\tilde{a}_i - \tilde{a}_j\|_1 \le \sqrt{K}\|\tilde{a}_i - \tilde{a}_j\|$. Moreover, since $\|\pi_i - \pi_j\|_\infty \le 1$, we have $\|\pi_i - \pi_j\| \le \|\pi_i - \pi_j\|_1$. It follows that

$$\|\pi_i - \pi_j\| \le C\|\tilde{a}_i - \tilde{a}_j\|. \tag{60}$$

This gives the second inequality in (58).

To get the first inequality in (58), introduce a vector $b \in \mathbb{R}^K$ with $b(k) = 1/V_1(k)$. Then (59) implies $\tilde{a}_j \propto (b \circ \pi_j)$ for all $1 \le j \le p$. Since both $\tilde{a}_j$ and π_j are weight vectors, we have $\tilde{a}_j = \frac{b \circ \pi_j}{\|b \circ \pi_j\|_1}$. Note that $C^{-1} \le \min_k V_1(k) \le \max_k V_1(k) \le C$ implies $C^{-1} \le \min_k b(k) \le \max_k b(k) \le C$. By replacing V_1 with b in the proof of (60), we immediately obtain

$$\|\tilde{a}_i - \tilde{a}_j\| \le C\|\pi_i - \pi_j\|.$$

This gives the second inequality in (58). □

7.3. *Proof of Lemma 6.3.* Introduce a set of p-dimensional random vectors $\{T_{im} : 1 \le i \le n, 1 \le m \le N\}$ such that they are independent of each other and that $T_{im} \sim \text{Multinomial}(1, d_i^0)$. From the model (4) and the definition of multinomial distributions,

$$z_i \overset{(d)}{=} \frac{1}{N} \sum_{m=1}^{N} (T_{im} - E[T_{im}]), \qquad 1 \le i \le n. \tag{61}$$

It follows that

$$M(j,j) - M_0(j,j) = \frac{1}{n} \sum_{i=1}^{n} z_i(j) \overset{(d)}{=} \frac{1}{Nn} \sum_{i=1}^{n} \sum_{m=1}^{N} \{T_{im}(j) - E[T_{im}(j)]\}.$$

Fix j and write $X_{im} = T_{im}(j) - E[T_{im}(j)]$. Then, $\{X_{im} : 1 \le i \le n, 1 \le m \le N\}$ are independent of each other. Moreover, since $T_{im}(j) \sim \text{Bernoulli}(d_i^0(j))$, we have $|X_{im}| \le 2$ and $\text{Var}(X_{im}) \le d_i^0(j) = \sum_{k=1}^{K} A_k(j) w_i(k) \le \sum_{k=1}^{K} A_k(j) = h_j$. We now apply the Bernstein inequality:

LEMMA 7.1 (Bernstein inequality). *Suppose $X_1, \cdots, X_n$ are independent random variables such that $EX_i = 0$, $|X_i| \le b$ and $\text{Var}(X_i) \le \sigma_i^2$ for all i. Let $\sigma^2 = n^{-1} \sum_{i=1}^{n} \sigma_i^2$. Then, for any $t > 0$,*

$$P\Big(n^{-1} |\sum_{i=1}^{n} X_i| \ge t\Big) \le 2 \exp\left(-\frac{nt^2/2}{\sigma^2 + bt/3}\right).$$

Using Lemma 7.1, we obtain

$$P\big(|M(j,j) - M_0(j,j)| \ge t\big) \le 2 \exp\left(-\frac{Nnt^2/2}{h_j + 2t/3}\right).$$

Let $t = (Nn)^{-1/2} \sqrt{10 h_j \log(n)}$. Since $h_j \ge h_{\min} \gg (Nn)^{-1} \log(n)$, we have $t \ll h_j$; therefore, in the denominator of the exponent, the term h_j is dominating. It follows that, with probability $1 - o(n^{-4})$,

$$|M(j,j) - M_0(j,j)| \le (Nn)^{-1/2} \sqrt{10 h_j \log(n)}.$$

According to the probability union bound, the above holds simultaneously for all $1 \le j \le p$ with probability $1 - o(pn^{-4}) = 1 - o(n^{-3})$.[9] □

[9] We have assumed $n \ge \max\{N, p\}$ without loss of generality. If $n < \max\{N, p\}$, the result continues to hold with $\log(n)$ replaced by $\log(\max\{n, N, p\})$.

7.4. *Proof of Lemma 6.4.* Consider the first claim. Fix k. Let $\{T_{im} : 1 \le i \le n, 1 \le m \le N\}$ be as in (61). It follows that

$$Z_j' W_k = \sum_{i=1}^{n} z_i(j) w_i(k) \overset{(d)}{=} \frac{1}{Nn} \sum_{i=1}^{n} \sum_{m=1}^{N} n w_i(k) \{T_{im}(j) - E[T_{im}(j)]\}.$$

Write $X_{im} = n w_i(k)\{T_{im}(j) - E[T_{im}(j)]\}$. Since $T_{im}(j) \sim \text{Bernoulli}(d_i^0(j))$, we find that $\text{Var}(X_{im}) \le n^2 w_i^2(k) d_i^0(j) \le n^2 h_j$ and $|X_{im}| \le 2n w_i(k) \le 2n$. We now apply Lemma 7.1 with $\sigma^2 = n^2 h_j$ and $b = 2n$. It yields that

$$P(|Z_j' W_k| > t) \le 2 \exp\left(\frac{Nnt^2/2}{n^2 h_j + 2nt/3} \right).$$

Set $t = C\sqrt{N^{-1} n h_j \log(n)}$ for a constant $C > 0$ to be decided. For such t, since $h_j \ge h_{\min} \gg (Nn)^{-1} \log(n)$, the term $n^2 h_j$ is the dominating term in the denominator of the exponent. Therefore, when C is properly large, the right hand side is $o(n^{-4})$. In other words, with probability $1 - o(n^{-4})$,

$$|Z_j' W_k| \le C N^{-1/2} \sqrt{n h_j \log(n)}. \tag{62}$$

Combing this with the probability union bound gives the claim.

Consider the second claim. Write

$$\|M_0^{-1/2} Z W_k\|^2 = \sum_{j=1}^{p} \frac{1}{M_0(j,j)} |Z_j' W_k|^2.$$

We have obtained the upper bound (62), which holds simultaneously for all $1 \le j \le p$, with probability $1 - o(n^{-3})$. Moreover, from (49), $M_0(j,j) \ge c_1 h_j$. As a result, with probability $1 - o(n^{-3})$,

$$\|M_0^{-1/2} Z W_k\|^2 \le \sum_{j=1}^{p} \frac{1}{c_1 h_j} \frac{C n h_j \log(n)}{N} = \frac{C n p \log(n)}{c_1 N}.$$

This proves the claim. □

7.5. *Proof of Lemma 6.5.* We aim to show that, for any given $1 \le j, \ell \le p$, with probability $1 - o(n^{-5})$,

$$\frac{1}{\sqrt{h_j h_\ell}} |Z_j' Z_\ell - E[Z_j' Z_\ell]| \le C N^{-1} \sqrt{n \log(n)}. \tag{63}$$

Once (63) is true, the claim follows from the probability union bound.

Below, we show (63). Fix (j,ℓ). Write $Z = [z_1,\ldots,z_n]$, and let $H = \mathrm{diag}(h_1,\ldots,h_p)$. Using the equality $xy = \frac{1}{4}(x+y)^2 - \frac{1}{4}(x-y)^2$, we find that

$$\begin{aligned}\frac{Z_j'Z_\ell}{\sqrt{h_jh_\ell}} &= \sum_{i=1}^n \frac{z_i(j)}{\sqrt{h_j}}\cdot\frac{z_i(\ell)}{\sqrt{h_\ell}}\\ &= \sum_{i=1}^n\left(\frac{z_i(j)}{2\sqrt{h_j}} + \frac{z_i(\ell)}{2\sqrt{h_\ell}}\right)^2 - \sum_{i=1}^n\left(\frac{z_i(j)}{2\sqrt{h_j}} - \frac{z_i(\ell)}{\sqrt{2h_\ell}}\right)^2\\ &= \sum_{i=1}^n (u_1'H^{-1/2}z_i)^2 - \sum_{i=1}^n (u_2'H^{-1/2}z_i)^2, \qquad u_1 \equiv \frac{e_j+e_\ell}{2}, u_2 \equiv \frac{e_j-e_\ell}{2};\end{aligned}$$

here $e_1,\ldots,e_p$ denote the standard basis vectors of $\mathbb{R}^p$. Taking the expectation on both sides, we find that $E[Z_j'Z_\ell]$ has a similar decomposition. As a result,

$$\begin{aligned}\frac{Z_j'Z_\ell - E[Z_j'Z_\ell]}{\sqrt{h_jh_\ell}} &= \sum_{i=1}^n\big\{(u_1'H^{-1/2}z_i)^2 - E[(u_1'H^{-1/2}z_i)^2]\big\}\\ &\quad - \sum_{i=1}^n\big\{(u_2'H^{-1/2}z_i)^2 - E[(u_2'H^{-1/2}z_i)^2]\big\}\\ &\equiv I + II. \qquad (64)\end{aligned}$$

First, we bound I. We start from studying $u_1'H^{-1/2}z_i$. Let $\{T_{im} : 1\le i\le n, 1\le m\le N\}$ be the same as in (61). It follows that

$$u_1'H^{-1/2}z_i \overset{(d)}{=} \frac{1}{N}\sum_{m=1}^N u_1'H^{-1/2}(T_{im} - E[T_{im}]).$$

Write $Y_{im} = u_1'H^{-1/2}(T_{im}-E[T_{im}])$. Since $T_{im}\sim\mathrm{Multinomial}(1,d_i^0)$, the covariance matrix of T_{im} equals to $\mathrm{diag}(d_i^0) - d_i^0(d_i^0)'$. It follows that $\mathrm{Var}(Y_{im}) \le u_1'H^{-1/2}\mathrm{diag}(d_i^0)H^{-1/2}u_1 = \frac{1}{4}(\frac{\sqrt{d_i^0(j)}}{\sqrt{h_j}} + \frac{\sqrt{d_i^0(\ell)}}{\sqrt{h_\ell}})^2 \le 1$, where the last inequality is because $d_i^0(j)\le h_j$. Furthermore, $|Y_{im}| \le 1/\sqrt{h_j} + 1/\sqrt{h_\ell} \le 2/\sqrt{h_{\min}}$. We now apply the Bernstein inequality, Lemma 7.1, with $\sigma^2 = 1$, $b = 2/\sqrt{h_{\min}}$. It gives

$$P\big(|u_1'H^{-1/2}z_i| > t\big) \le 2\exp\left(-\frac{Nt^2/2}{1+2t/(3\sqrt{h_{\min}})}\right), \qquad \text{for all } t>0. \qquad (65)$$

Let $t_0 = CN^{-1/2}\sqrt{\log(n)}$ for a large enough $C>0$. Since $Nh_{\min}\gg\log(n)$, for all $0<t\le t_0$, the right hand side of (65) is bounded by $2e^{-Nt^2/4}$. Define

$$X_i = (u_1'H^{-1/2}z_i)\cdot 1\big\{|u_1'H^{-1/2}z_i|\le t_0\big\}.$$

When C is chosen properly large, we have the following results:

(i) $X_i = u_1' H^{-1/2} z_i$ with probability $1 - o(n^{-6})$.
(ii) X_i is a sub-Gaussian random variable with the sub-Gaussian norm $\|X_i\|_{\psi_2} = O(1/\sqrt{N})$.
(iii) $|E[(u'H^{-1/2}z_i)^2] - E[X_i^2]| = o(n^{-5})$.

Here (i) is because $P(X_i \neq u_1' H^{-1/2} z_i) = P(|u_1' H^{-1/2} z_i| > t_0) \leq 2e^{-Nt_0^2/4} = O(n^{-C^2/4})$; (ii) is because: for $0 < t \leq t_0$, $P(|X_i| > t) \leq P(|u_1' H^{-1/2} z_i| > t) \leq 2e^{-Nt^2/4}$, and for $t > t_0$, $P(|X_i| > t) = 0$; (iii) is because $|E[(u'H^{-1/2}z_i)^2] - E[X_i^2]| \leq (2/\sqrt{h_{\min}})^2 \cdot P(|u'H^{-1/2}z_i| > t_0) = o(N) \cdot O(n^{-C^2/4})$. Using (i)-(iii) above, with probability $1 - o(n^{-5})$,

$$I = \sum_{i=1}^{n} (X_i^2 - E[(u_1' H^{-1/2} z_i)]) = \sum_{i=1}^{n} (X_i^2 - E[X_i^2]) + o(n^{-4}). \tag{66}$$

Since each X_i is sub-Gaussian, $X_i^2 - E[X_i^2]$ is a sub-exponential random variable with the sub-exponential norm $\|X_i^2 - E[X_i^2]\|_{\psi_1} \leq 2\|X_i\|_{\psi_2}^2 = O(1/N)$ (Vershynin, 2012, Lemma 5.14, Remark 5.18). We apply the Bernstein's inequality for sub-exponential variables (Vershynin, 2012, Corollary 5.17):

LEMMA 7.2 (Bernstein's inequality for sub-exponential variables). *Suppose $X_1, \cdots, X_n$ are independent random variables such that $EX_i = 0$ and $\max_{1 \leq i \leq n} \|X\|_{\psi_1} \leq \kappa$. Then, for any $t > 0$,*

$$P\Big(|\sum_{i=1}^{n} X_i| > nt\Big) \leq 2\exp\left(-cn \min\left\{\frac{t^2}{\kappa^2}, \frac{t}{\kappa}\right\}\right),$$

where $c > 0$ is a universal constant.

We apply Lemma 7.2 with $\kappa = C_1/N$ and $t = C_2\kappa\sqrt{n^{-1}\log(n)}$ for $C_1, C_2 > 0$ that are large enough. It follows that with probability $1 - o(n^{-5})$, $|\sum_{i=1}^{n}(X_i^2 - E[X_i^2])| \leq CN^{-1}\sqrt{n\log(n)}$. Combining it with (66) gives: with probability $1 - o(n^{-5})$,

$$|I| \leq CN^{-1}\sqrt{n\log(n)}. \tag{67}$$

Next, we bound II. When $j = \ell$, II is exactly equal to 0. When $j \neq \ell$, we can similarly write $u_2' H^{-1/2} z_i = N^{-1} \sum_{m=1}^{N} Y_{im}$, with $Y_{im} = u_2' H^{-1/2}(T_{im} - E[T_{im}])$. Then, $|Y_{im}| \leq \max\{1/\sqrt{h_j}, 1/\sqrt{h_\ell}\} \leq 1/\sqrt{h_{\min}}$, and $\mathrm{Var}(Y_{im}) \leq u_2' H^{-1} \mathrm{diag}(d_i^0) H^{-1/2} u_2 \leq \frac{1}{4}(\frac{\sqrt{d_i^0(j)}}{\sqrt{h_j}} - \frac{\sqrt{d_i^0(\ell)}}{\sqrt{h_\ell}})^2 \leq \frac{1}{4}$. We again apply Lemma 7.1

to bound the tail probability of $u_2'H^{-1/2}z_i$, and then apply Lemma 7.2 to bound II. Similarly, we find that, with probability $1-o(n^{-5})$,

$$|II| \le CN^{-1}\sqrt{n\log(n)}. \tag{68}$$

Then, (63) follows from plugging (67)-(68) into (64). □

7.6. *Proof of Lemma 6.6.* Let $H=\mathrm{diag}(h_1,\ldots,h_p)$. By (49), $M_0(j,j)\ge c_1h_j$ for all $1\le j\le p$. It follows that $\|M_0^{-1/2}H^{1/2}\|\le c_1^{-1/2}$. As a result,

$$\begin{aligned}&\|M_0^{-1/2}(ZZ'-E[ZZ'])M_0^{-1/2}\|\\ =&\|M_0^{-1/2}H^{1/2}\|\cdot\|H^{-1/2}(ZZ'-E[ZZ'])H^{-1/2}\|\cdot\|H^{1/2}M_0^{-1/2}\|\\ \le&c_1^{-1}\|H^{-1/2}(ZZ'-E[ZZ'])H^{-1/2}\|.\end{aligned}$$

Therefore, to show the claim, it suffices to show that

$$\|H^{-1/2}(ZZ'-E[ZZ'])H^{-1/2}\|\le C\Big(\frac{1}{N}+\frac{p}{N^2h_{\min}}\Big)\sqrt{np}. \tag{69}$$

To show (69), we need some existing results on α-nets. For any $\alpha>0$, a subset $\mathcal{M}$ of the unit sphere $\mathcal{S}^{p-1}$ is called an α-net if $\sup_{x\in\mathcal{S}^{p-1}}\inf_{y\in\mathcal{M}}\|x-y\|\le\alpha$. The following lemma combines Lemmas 5.2-5.3 in Vershynin (2012).

LEMMA 7.3 (α-net). *Fix $\alpha\in(0,1/2)$. There exists an α-net $\mathcal{M}_\alpha$ of $\mathcal{S}^{p-1}$ such that $|\mathcal{M}_\alpha|\le(1+2/\alpha)^p$. Moreover, for any symmetric $p\times p$ matrix B, $\|B\|\le(1-2\alpha)^{-1}\sup_{u\in\mathcal{M}_\alpha}\{|u'Bu|\}$.*

By Lemma 7.3, there exists a $(1/4)$-net $\mathcal{M}_{1/4}$, such that $|\mathcal{M}_{1/4}|\le 9^p$ and

$$\|H^{-1/2}(ZZ'-E[ZZ'])H^{-1/2}\|\le 2\max_{u\in\mathcal{M}_{1/4}}\{|u'H^{-1/2}(ZZ'-E[ZZ'])H^{-1/2}u|\}.$$

Therefore, to show (69), it is sufficient to show that, for any fixed $u\in\mathcal{S}^{p-1}$, with probability $1-o(9^{-p}n^{-3})$,

$$|u'H^{-1/2}(ZZ'-E[ZZ'])H^{-1/2}u|\le C\Big(\frac{1}{N}+\frac{p}{N^2h_{\min}}\Big)\sqrt{np}. \tag{70}$$

Below, we show (70). Write $Z=[z_1,\ldots,z_n]$. For any $u\in\mathcal{S}^{p-1}$,

$$\begin{aligned}&u'H^{-1/2}(ZZ'-E[ZZ'])H^{-1/2}u\\ =&\sum_{i=1}^n\{(u'H^{-1/2}z_i)^2-E[(u'H^{-1/2}z_i)^2]\}.\end{aligned} \tag{71}$$

Our plan is to first get a tail bound for $u'H^{-1/2}z_i$, which is similar to (65). We then consider two separate cases, $Nh_{\min} \geq p$ and $Nh_{\min} < p$: for each case, we use the tail bound of $u'H^{-1/2}z_i$ to prove (70).

First, we study $u'H^{-1/2}z_i$. Let $\{T_{im} : 1 \leq i \leq n, 1 \leq m \leq N\}$ be the set of random variables as in (61). Write

$$u'H^{-1/2}z_i \stackrel{(d)}{=} \frac{1}{N}\sum_{m=1}^{N} Y_{im}, \qquad \text{with } Y_{im} = u'H^{-1/2}(T_{im} - E[T_{im}]).$$

Since T_{im} follows a distribution of Multinomial$(1, d_i^0)$, it is easy to see that $|Y_{im}| \leq 2/\sqrt{h_{\min}}$ and $\mathrm{var}(Y_{im}) \leq u'H^{-1/2}\mathrm{diag}(d_i^0)H^{-1/2}u \leq \|u\|^2 \leq 1$ (note that $d_i^0(j) = \sum_{k=1}^K A_k(j)w_i(k) \leq \sum_{k=1}^K A_k(j) = h_j$). We apply the Bernstein's inequality, Lemma 7.1, and obtain that, for any $t > 0$,

$$P(|u'H^{-1/2}z_i| > t) \leq 2\exp\left(-\frac{Nt^2/2}{1 + 2t/(3\sqrt{h_{\min}})}\right), \qquad \text{for all } t > 0. \tag{72}$$

Next, we prove (70) for two cases separately: $Nh_{\min} \geq p$ and $Nh_{\min} < p$. In the first case, for a constant $C_1 > 0$ to be decided, let $\delta_{n1} = C_1\sqrt{p/N}$. Since $Nh_{\min} \geq p$, we have

$$P(|u'H^{-1/2}z_i| > t) \leq 2\exp\left(-\frac{Nt^2/2}{1 + 2C_1/3}\right), \qquad \text{for all } 0 < t \leq \delta_{n1}. \tag{73}$$

We then define a truncated version of $u'H^{-1/2}z_i$:

$$X_i \equiv u'H^{-1/2}z_i \cdot 1\{|u'H^{-1/2}z_i| \leq \delta_{n1}\}, \qquad 1 \leq i \leq n.$$

We claim that

(i) $X_i = u'H^{-1/2}z_i$ with probability $1 - o(9^{-p}n^{-4})$.
(ii) X_i is a sub-Gaussian random variable with the sub-Gaussian norm $\|X_i\|_{\psi_2} = O(1/\sqrt{N})$.
(iii) $|E[(u'H^{-1/2}z_i)^2] - E[X_i^2]|$ is negligible compared with the right hand side of (70).

Here (ii) is a direct result of (73). To see (i), note that by (73), $P(|u'H^{-1/2}z_i| > \delta_{n1}) \leq 2\exp(-\frac{C_1^2/2}{1+2C_1/3}p)$; since $p \gg \log(n)$, with an appropriately large C_1, this probability is $o(9^{-p}n^{-4})$. To see (iii), note that $|u'H^{-1/2}z_i| \leq 2/\sqrt{h_{\min}} \leq 2\sqrt{N/p}$; so, $|E[(u'H^{-1/2}z_i)^2] - E[X_i^2]| \leq (4N/p) \cdot P(|u'H^{-1/2}z_i| > \delta_{n1}) \leq (8N/p) \cdot \exp(-\frac{C_1^2/2}{1+2C_1/3}p)$. Since $p \gg \log(N+n)$, when C_1 is large enough,

this quantity is $o(N^{-1}\sqrt{np})$. Combining (i)-(iii) with (71), with probability $1-o(9^{-p}n^{-3})$,

$$|u'H^{-1/2}(ZZ'-E[ZZ'])H^{-1/2}u| \le |\sum_{i=1}^{n}(X_i^2-E[X_i^2])| + o(N^{-1}\sqrt{np}). \tag{74}$$

Since each X_i is sub-Gaussian, $X_i^2-E[X_i^2]$ is a sub-exponential random variable with the sub-exponential norm $\|X_i^2-E[X_i^2]\|_{\psi_1} \le 2\|X_i\|_{\psi_2}^2 = O(1/N)$ (Vershynin, 2012, Lemma 5.14, Remark 5.18). We then apply Lemma 7.2 with $\kappa = O(1/N)$ and $t = C\kappa\cdot\sqrt{p/n}$. When the constant C is large enough, with probability $1-o(9^{-p}n^{-3})$,

$$|\sum_{i=1}^{n}(X_i^2-E[X_i^2])| \le nt \le CN^{-1}\sqrt{np}. \tag{75}$$

Combining (74)-(75) gives (70) in the first case.

In the second case, let $\delta_{n2} = C_2 p/(N\sqrt{h_{\min}})$ for a constant $C_2 > 0$ to be determined. We study the right hand of (72). Note that $Nh_{\min} < p$. For $t \le \delta_{n2}$, we have $1+2t/(3\sqrt{h_{\min}}) \le p/(Nh_{\min}) + 2\delta_{n2}/(3\sqrt{h_{\min}}) = (1+2C_2/3)\cdot p/(Nh_{\min})$; for $t > \delta_{n2}$, we have $1+2t/(3\sqrt{h_{\min}}) \le \delta_{n2}/(C_2\sqrt{h_{\min}}) + 2t/(3\sqrt{h_{\min}}) = (C_2^{-1}+2/3)\cdot t/\sqrt{h_{\min}}$. Plugging them into (72) gives

$$P(|u'H^{-1/2}z_i| > t) \le 2\begin{cases} \exp\Big(-\frac{1/2}{1+2C_2/3}\cdot p^{-1}N^2h_{\min}\cdot t^2\Big), & \text{for } 0 < t \le \delta_{n2}, \\ \exp\Big(-\frac{1/2}{C_2^{-1}+2/3}\cdot N\sqrt{h_{\min}}\cdot t\Big), & \text{for } t > \delta_{n2}. \end{cases} \tag{76}$$

In particular, $P(|u'H^{-1/2}z_i| > \delta_{n2}) \le 2e^{-\frac{3C_2^2}{6+4C_2}p}$. In light of this, we introduce a truncated version of $u'H^{-1/2}z_i$:

$$\tilde{X}_i \equiv u'H^{-1/2}z_i\cdot 1\{|u'H^{-1/2}z_i| \le \delta_{n2}\}, \qquad 1 \le i \le n.$$

We have the following observations, whose proofs are similar to the (i)-(iii) in the first case and are omitted.

(i) $\tilde{X}_i = u'H^{-1/2}z_i$ with probability $1-o(9^{-p}n^{-4})$.
(ii) $\tilde{X}_i$ is a sub-Gaussian random variable with the sub-Gaussian norm $\|\tilde{X}_i\|_{\psi_2} = O(\sqrt{p/(N^2h_{\min})})$.
(iii) $|E[(u'H^{-1/2}z_i)^2] - E[\tilde{X}_i^2]|$ is negligible compared with the right hand side of (70).

From (ii), $\tilde{X}_i^2 - E[\tilde{X}_i^2]$ is a sub-exponential random variable with the sub-exponential norm $\|\tilde{X}_i^2 - E[\tilde{X}_i^2]\|_{\psi_1} = O(p/(N^2h_{\min}))$. We apply Lemma 7.2

with $\kappa = O(p/(N^2 h_{\min}))$ and $t = O(\kappa\sqrt{p/n})$. Combining the result with (i) and (iii), we find that, with probability $1 - o(9^{-p}n^{-3})$,

$$|u'H^{-1/2}(ZZ' - E[ZZ'])H^{-1/2}u| \le |\sum_{i=1}^{n}(\tilde{X}_i^2 - E[\tilde{X}_i^2])| + o\big(\frac{p\sqrt{np}}{N^2 h_{\min}}\big)$$

$$\le Cn\kappa\sqrt{p/n} + o\big(\frac{p\sqrt{np}}{N^2 h_{\min}}\big) \le \frac{Cp\sqrt{np}}{N^2 h_{\min}}. \tag{77}$$

This proves (70) in the second case. □

Acknowledgements. The authors would like to thank Jiashun Jin, John Lafferty, and Art Owen for helpful pointers.

References.

Airoldi, E., Blei, D., Fienberg, S. and Xing, E. (2008). Mixed membership stochastic blockmodels. *J. Mach. Learn. Res.* **9** 1981–2014.

Arias-Castro, E., Lerman, G. and Zhang, T. (2017). Spectral clustering based on local PCA. *J. Mach. Learn. Res.* **18** 1–57.

Arora, S., Ge, R. and Moitra, A. (2012). Learning topic models–going beyond SVD. In *Foundations of Computer Science (FOCS)* 1–10.

Arora, S., Ge, R., Halpern, Y., Mimno, D., Moitra, A., Sontag, D., Wu, Y. and Zhu, M. (2013). A practical algorithm for topic modeling with provable guarantees. In *International Conference on Machine Learning (ICML)* 280–288.

Bansal, T., Bhattacharyya, C. and Kannan, R. (2014). A provable SVD-based algorithm for learning topics in dominant admixture corpus. In *Adv. Neural Inf. Process. Syst.* 1997–2005.

Berthet, Q. and Rigollet, P. (2013). Optimal detection of sparse principal components in high dimension. *Ann. Statist.* **41** 1780–1815.

Birnbaum, A., Johnstone, I. M., Nadler, B. and Paul, D. (2013). Minimax bounds for sparse PCA with noisy high-dimensional data. *Ann. Statist.* **41** 1055.

Blei, D. (2012). Probabilistic topic models. *Commun. ACM* **55** 77–84.

Blei, D. and Lafferty, J. (2007). A correlated topic model of science. *Ann. Appl. Statist.* **1** 17–35.

Blei, D., Ng, A. and Jordan, M. (2003). Latent dirichlet allocation. *J. Mach. Learn. Res.* **3** 993–1022.

Davis, C. and Kahan, W. M. (1970). The rotation of eigenvectors by a perturbation. III. *SIAM J. Numer. Anal.* **7** 1–46.

Ding, W., Rohban, M. H., Ishwar, P. and Saligrama, V. (2013). Topic discovery through data dependent and random projections. In *International Conference on Machine Learning (ICML)* 1202–1210.

Donoho, D. and Stodden, V. (2004). When does non-negative matrix factorization give a correct decomposition into parts? In *Adv. Neural Inf. Process. Syst.* 1141–1148.

Fan, J., Fan, Y. and Lv, J. (2008). High dimensional covariance matrix estimation using a factor model. *J. Econometrics* **147** 186–197.

Fan, J., Liao, Y. and Mincheva, M. (2011). High dimensional covariance matrix estimation in approximate factor models. *Ann. Statist.* **39** 3320.

Han, F. and Liu, H. (2014). Scale-invariant sparse PCA on high-dimensional meta-elliptical data. *J. Amer. Statist. Soc.* **109** 275–287.

HARMAN, D. (1993). Overview of the first Text REtrieval Conference (TREC-1). In *Proceedings of the first Text REtrieval Conference (TREC-1)* 1–20.

HOFMANN, T. (1999). Probabilistic latent semantic indexing. In *International ACM SIGIR conference* 50–57.

HORN, R. and JOHNSON, C. (1985). *Matrix Analysis*. Cambridge University Press.

JI, P. and JIN, J. (2016). Coauthorship and citation networks for statisticians. *Ann. Appl. Statist.* **10** 1779–1812.

JIN, J. (2015). Fast community detection by SCORE. *Ann. Statist.* **43** 57–89.

JIN, J., KE, Z. T. and LUO, S. (2016). Estimating network memberships by simplex vertices hunting. *Manuscript*.

JIN, J. and WANG, W. (2016). Influential Features PCA for high dimensional clustering. *Ann. Statist.* **44** 2323–2359.

JOHNSTONE, I. and LU, A. (2009). On consistency and sparsity for principal components analysis in high dimensions. *J. Amer. Statist. Assoc.* **104**.

KE, Z. T. and WANG, M. (2017). A spectral approach to nonnegative matrix factorization. *Manuscript*.

KOLAR, M. and TADDY, M. Discussion of "Coauthorship and citation networks for statisticians". *Ann. Appl. Stat.* **4** 1835–1841.

LEE, D. and SEUNG, S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature* **401** 788–791.

OWEN, A. and WANG, J. (2016). Bi-cross-validation for factor analysis. *Statist. Sci.* **31** 119–139.

PAATERO, P. and TAPPER, U. (1994). Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics* **5** 111–126.

PAPADIMITRIOU, C. H., RAGHAVAN, P., TAMAKI, H. and VEMPALA, S. (2000). Latent semantic indexing: A probabilistic analysis. *J. Comput. System Sci.* **61** 217 - 235.

PERRONE, V., JENKINS, P. A., SPANO, D. and TEH, Y. W. (2016). Poisson Random Fields for Dynamic Feature Models. *arXiv:1611.07460*.

SALDANA, D., YU, Y. and FENG, Y. (2017). How many communities are there? *J. Comp. Graph. Stat.* **26** 171–181.

TSYBAKOV, A. B. (2009). Introduction to nonparametric estimation. Revised and extended from the 2004 French original. Translated by Vladimir Zaiats.

VERSHYNIN, R. (2012). Introduction to the non-asymptotic analysis of random matrices. In *Compressed Sensing: Theory and Applications* 210–268. Cambridge Univ. Press.

VU, V. Q. and LEI, J. (2013). Minimax sparse principal subspace estimation in high dimensions. *Ann. Statist.* **41** 2905–2947.

WANG, T., BERTHET, Q. and SAMWORTH, R. J. (2016). Statistical and computational trade-offs in estimation of sparse principal components. *Ann. Statist.* **44** 1896–1930.

ZOU, H., HASTIE, T. and TIBSHIRANI, R. (2006). Sparse principal component analysis. *J. Comp. Graph. Stat.* **15** 265–286.

Z. KE AND M. WANG
DEPARTMENT OF STATISTICS
UNIVERSITY OF CHICAGO
CHICAGO, ILLINOIS, IL 60637
E-MAIL: zke@galton.uchicago.edu
minzhew@galton.uchicago.edu